AD-A036 069 AUERBACH ASSOCIATES INC PHILADELPHIA PA F/G 9/2

THE ANALYSIS OF CREDIBILITY AND CONSISTENCY IN INTELLIGENCE DAT--ETC(U)

DEC 76 J SABLE, R DICKSON F30602-75-C-0330

UNCLASSIFIED AAI-2329-TR-1 RADC-TR-76-392 NL

1 OF 3

AD A036069

12

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

19 **REPORT DOCUMENTATION PAGE**

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER 18 RADC-TR-76-392	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) 6 THE ANALYSIS OF CREDIBILITY AND CONSISTENCY IN INTELLIGENCE DATA.		5. TYPE OF REPORT & PERIOD COVERED 9 Final Technical Report, Jun 75 - Sep 76 6. PERFORMING ORG. REPORT NUMBER 14 AAI-2329-TR-1
7. AUTHOR(s) 10 Jerome Sable Robert Dickson		8. CONTRACT OR GRANT NUMBER(s) 15 F30602-75-C-0330
9. PERFORMING ORGANIZATION NAME AND ADDRESS Auerbach Associates Inc 121 North Broad Street Philadelphia PA 19107		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 64750F 20530202
11. CONTROLLING OFFICE NAME AND ADDRESS Rome Air Development Center (IRDT) Griffiss AFB NY 13441		12. REPORT DATE 11 December 1976 13. NUMBER OF PAGES 204
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Same 12 211p.		15. SECURITY CLASS. (of this report) UNCLASSIFIED 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE N/A

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

Same

16 2053 17 02

18. SUPPLEMENTARY NOTES

RADC Project Engineer:
Robert N. Ruberti (IRDT)

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Intelligence Data Handling	Fuzzy Logic
Data Credibility Analysis	LISP Modifications
Data Consistency Analysis	
Probability Theory	

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report describes methods for the analysis, quantification and calculation of intelligence data credibility, methods for analyzing the consistency of facts stored in a scientific intelligence data base, and modifications to the Univac 1110 LISP programming language required for implementation of a scientific intelligence information system.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

391061

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

PREFACE

The work reported here was accomplished under Contract Number F30602-75-C-0330 under sponsorship of the Rome Air Development Center (RADC). The principal objective of this project was to develop methods for handling credibility and consistency aspects of data in an intelligence system with advanced inference capabilities. It is directed in particular to the Scientific and Technical Intelligence System (STIS) being developed at the Air Force Foreign Technology Division (FTD). The report was prepared by Robert Dickson and Jerome Sable of AAI (AUERBACH Associates, Inc.). Mr. Ken Rose was also part of the project team and developed theoretical aspects of this work (e.g., Appendix B). Grateful acknowledgement is given to Edward Stull of FTD and Robert Ruberti of RADC. Dr. Sable was the Project Manager.

In a separate task ammended to this contract AAI has modified the LISP system for the UNIVAC 1110 Computer to make it more suitable for implementation of new versions of STIS. This work is also reported in this volume.

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

TABLE OF CONTENTS (Continued)

TABLE OF CONTENTS (Continued)

LIST OF ILLUSTRATIONS

LIST OF TABLES

EVALUATION

Functional design specifications for methods of handling the credibility and consistency of facts in an intelligence data base have been delivered under Contract F30602-75-C-0330. When implemented as an adjunct to an intelligence information system, these methods will enhance analyst-inferential capabilities. Future application of these methods is planned for the Scientific and Technical Information System (STIS) at the Foreign Technology Division. This development is included as part of TPO No 3, Indications and Warning.

Robert N Ruberti

ROBERT N. RUBERTI
Project Engineer

(The reverse of this page is blank)

SECTION I. INTRODUCTION

The intelligence analyst must often cope with large volumes of information which he uses as clues to construct a true representation of the state of affairs in the real world. As he attempts to build a description of the entities which he is interested in, the analyst often finds that they are richly interrelated and that the description of a typical entity is quite fragmentary. Compounding his problem is the fact that the entities he is attempting to describe are usually not perceived directly but are known only through reports and sensors with various levels of credibility and accuracy.

Under sponsorship of Rome Air Development Center, AUERBACH Associates, Inc. (AAI) has been studying the problem of intelligence data processing, and developing advanced data structures and inference techniques, and assisting the Air Force Foreign Technology Division (FTD) in developing an advanced intelligence system called STIS (Scientific and Technical Intelligence System). Allowing the analyst to enter general rules and credibility judgements, in addition to explicit facts, and providing a capability to derive implicit results, raises new questions concerning information credibility and consistency which are addressed in this report.

1.1 OBJECTIVES

The objectives of this project were to:

(a) study and develop a methodology for representing the credibility of information in the STIS data base and the results derived from it,

(b) develop a methodology for determining the logical consistency of delineated subsets of facts and rules in the STIS data base,

(c) augment the LISP system for the UNIVAC 1100 series computer to provide a more suitable implementation language for advanced STIS capabilities.

This work was carried out in the context of STIS as it is being developed through the effort of FTD, with support from AAI. In particular, this work is consistent with STIS information structures and inference strategies.

1.2 BACKGROUND

STIS provides an advanced capability for the analysis of intelligence information. It is based on a network type data structure which permits relationships among entities and new attributes to be freely defined with minimal impact on previously stored data and programs. Because of this, it is particularly suited for capturing fragmentary information which is undergoing collation processing, analysis, evaluation, and synthesis into finished intelligence. A description of the STIS information structure, called the Concept Net, is given in Appendix A to this report.

In another effort, AAI developed the design of advanced relational data and inference providing tools for use in an operational intelligence environment. It is expected that STIS will be used as a vehicle with which to develop and test operations on relations, inference, and consistency determining functions. The ultimate goal is to incorporate these advanced capabilities into STIS so that their effectiveness can be accurately evaluated, and these new tools can be provided to the STIS analyst.

AAI has developed the design of an inference capability for STIS at FTD. An initial approach to the derivation of bounds on the truth values of derived results was developed (Sable) and extensions to that method are described in this report. Clearly, the ultimate credibility and consistency scheme will have to be closely integrated with the inference capability and other STIS mechanisms currently being developed.

1.3 ROLES OF CREDIBILITY AND CONSISTENCY IN INTELLIGENCE ANALYSIS

The art of intelligence, by its very nature, is intimately concerned with data of variable and sometimes unknown accuracy. The analyst is confronted with a two-stage problem. First, he must discern, from the data at hand, what is the most likely state-of-affairs in the real world, and how accurate and complete that picture is apt to be. Second, he must use the inferential tools at his disposal, both formal (deductive and inductive) and intuitive, to estimate the implications of the current state-of-affairs relative to particular questions facing him. The questions, often of a composite nature and not immediately evident from the current state-of-affairs, may be one or more of the following:

- Does the current state-of-affairs represent a significant (unexpected) departure from a prior state-of-affairs?
- Does it represent a significantly new level of enemy capability?
- What development directions are indicated and further capabilities implied?
- How does it compare with our own capabilities?
- What does it imply about the enemy's plans and motives?
- Does it represent a threat (technological or otherwise) to ourselves?

Clearly, these questions cannot be answered automatically with today's information technology. The objective of long term Scientific and Technical Intelligence System (STIS) developments is to furnish the analyst with effective tools with

which to build and maintain a model of the current state-of-affairs which is as realistic and complete as possible. It should be able to represent "facts" (with his judgement of their credibility) in terms which have meaning to him, and to represent plausible rules of inference and their level of credibility. An inferential system is being designed which may make it possible to use these rules to derive implicit facts together with a bound on their credibility. The facts, rules, and steps taken in the inferential process will be displayed to the analyst so that he can judge the validity of the process and the credibility of the result.

What is meant by the intuitive notion of credibility is itself a complex question which deserves careful development. One can view the intelligence observation and analysis process as including the elements shown in Figure 1-1. This is the event input chain which produces "facts" for the data base. The analyst, besides entering these facts, must judge their credibility relative to the current state-of-affairs and their value relative to current and anticipated questions of the type listed above. This process involves a number of factors including the analyst's evaluation of the source and sensor, the source's evaluation of the sensor, and the quality of the observation. Some of these evaluations by the originator (source) and destination analyst are shown in Table 1-1. An indication of how these and other factors relate to the overall evaluation of credibility and utility (value) by the analyst is given in the Factor Dependence Tree of Figure 1-2. This diagram attempts to show how the worth of a fact or report can be considered as being influenced by the factors of relevance, credibility, and analysis or assimilation level and how credibility and other factors are, in turn, influenced by more basic factors such as source reliability, accuracy, consistency, etc. Clearly, it is unrealistic to believe that the credibility evaluation process can be completely formalized. A less naive view is that some of the important factors can be quantified and enough information made available to the analyst so that valid judgements can be made.

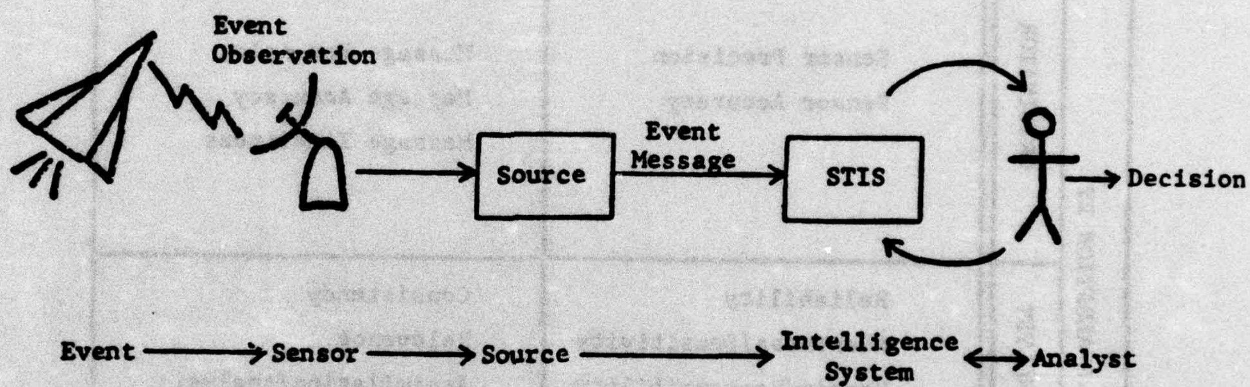

Figure 1-1. Event Input Chain

TABLE 1-1. INFORMATION ASSESSMENT FACTORS

EVALUATION BY		EVALUATION OF: SENSOR OR SOURCE	EVALUATION OF: FACT, MESSAGE, OR RESULT
EVALUATION BY	ORIGINATOR	Sensor Precision Sensor Accuracy	Message Precision Message Accuracy Message Timeliness
	DESTINATION ANALYST	Reliability Alertness/Sensitivity Domain/Responsibility Sensor Capability Source Bias	Consistency Relevance Assimilation/Analysis Level Credibility Value

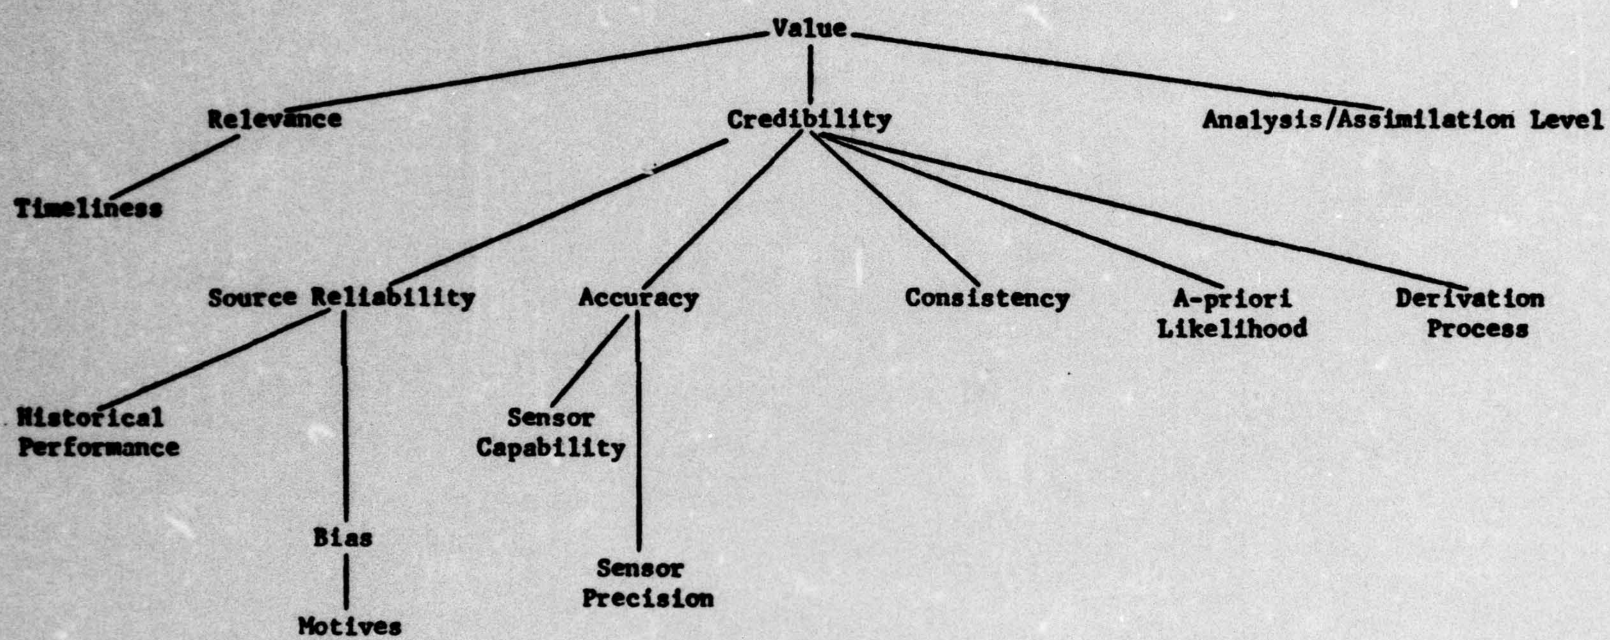

Figure 1-2. Factor Dependence Tree

(The reverse of this page is blank)

SECTION II. SUMMARY

This project consisted of three largely independent areas:

(1) The analysis, quantification, and calculation of credibility,

(2) The analysis of fact/rule consistency in a technical intelligence Concept Net, and

(3) Modification of the LISP programming system for the UNIVAC 1110 to make it more suitable for implementation of portions of STIS.

The first two areas involve the specification of procedures which will contribute inherent intelligence analysis capabilities to STIS. The third provides an additional systems implementation tool for the UNIVAC 1110.

2.1 CREDIBILITY ANALYSIS

2.1.1 Mathematical Foundations For Handling Credibility

The analysis of credibility of both explicit and derived information is a complex question that may be approached from several points of view, utilizing at least the mathematical models of probability theory and multi-valued logic. As we shall see, "fuzzy" logic (Lee) is a special case of multi-valued logics. We examine first probability theory, then multi-valued logic as a theoretical foundation (mathematical model) of credibility.

2.1.1.1 Probability Theory

Credibility, or the truth value of a sentence, can be viewed as an interpretation of probability. This interpretation of probability is sometimes called inductive logic (Carnap and Jeffrey), the logic of weight (Reichenbach), or more commonly subjective probability. It has been applied to the problem of determining the probability of an event, given reports of the event of various reliabilities. This problem has been examined from an intelligence analysis point of view by (Kuhns) and by (Johnscn). It has also been examined, in a more general scientific and behavioral framework by (Shum), (Cavanagh), and (Snapper).

Often the task is formally analogous to a problem in statistical inference, where items of evidence or data are used to determine the relative likelihood of alternative hypotheses. One strategy for processing data in two tasks is Bayes' theorem, a form of which is : *

$$\langle D \rightarrow Hi \rangle = \frac{\langle Hi \rightarrow D \rangle \langle Hi \rangle}{\sum_{i} \langle Hi \rightarrow D \rangle \langle Hi \rangle}$$

---

* For conciseness, and to preserve a parallel with deductive logic, we use the following notation in this report:

$\langle a \rangle$ is the probability that event a occurs (is true), usually written P(a).

$\langle \bar{a} \rangle$ or $\langle \neg a \rangle$ is the probability that event a does not occur (is false), $\langle \bar{a} \rangle = 1 - \langle a \rangle$.

$\langle a \rightarrow b \rangle$ is the conditional probability from a to b, the probability that b occurs given event a, usually written $P(b|a)$.

where $\langle Hi \rangle$ is the prior probability of a particular hypothesis; $\langle Hi \rightarrow D \rangle$ is the probability of the occurrence of a particular item of data conditional upon the truth of a particular hypothesis; and $\langle D \rightarrow Hi \rangle$ is the posterior probability of a particular hypothesis conditional upon the occurrence of a particular datum. Expressed in this way, the estimation of posterior probability is seen to involve two processes: first, the determination of the diagnostic impact of each datum ($\langle Hi \rightarrow D \rangle$); and second, calculation of the posterior probability estimate ($\langle D \rightarrow Hi \rangle$) on the basis of the observed data.

In the long term STIS design both specific information and general rules are stored in a relational network called a Concept Net. Both the specific information (such as "facts" describing entities and their interrelationships) and the rules are associated with a probabilistic measure of their truth value called credibility. Rules are entered as logical implications of the form $e \Rightarrow h$ where e and h may be logical propositions. Typically e is a conjunction of terms (relations) and h is a single derived relation. The credibility, or strength, of rules is given by a pair of numbers which represent the conditional probabilities $\{\langle e \rightarrow h \rangle, \langle \bar{e} \rightarrow h \rangle\}$. Given the existence (or derivability) of the premise e with credibility $\langle e \rangle$, the conclusion h can be derived with credibility $\langle h \rangle$ using the following derivation.

$$\langle h \rangle = \langle e \wedge h \rangle + \langle \bar{e} \wedge h \rangle = \langle e \rangle \cdot \langle e \rightarrow h \rangle + \langle \bar{e} \rangle \cdot \langle \bar{e} \rightarrow h \rangle$$

The credibility $\langle e \rangle$ represents the subjective (prior) probability that the premise is true. If e is a compound proposition, its credibility can be derived from its components using the laws of probability. For example if $e = e_1 \wedge e_2$ then $\langle e \rangle = \langle e_1 \rangle \cdot \langle e_1 \rightarrow e_2 \rangle = \langle e_2 \rangle \cdot \langle e_2 \rightarrow e_1 \rangle$. If the terms e_1 and e_2 are independent then $\langle e \rangle = \langle e_1 \rangle \cdot \langle e_2 \rangle$. The rule strength factor $\langle e \rightarrow h \rangle$ is the probability that the conclusion (hypothesis) is true given that the premise (evidence) is true. The rule strength factor $\langle \bar{e} \rightarrow h \rangle$ is the probability that the conclusion is true even when the premise is false. A rule, which produces a hypothesis with a high credibility may have a high value not only for $\langle e \rightarrow h \rangle$. However, an effective rule, one which yields an hypothesis of high credibility only when the evidence has high credibility, should have a high likelihood ratio $\langle e \rightarrow h \rangle / \langle \bar{e} \rightarrow h \rangle$.

The credibility of the evidence is enhanced when independent reports of the same event are received. If we let E_i represent a report of event e from source S_i then the conditional probability $<E_i \rightarrow e>$ represents the credibility of e due to the single report. If two independent sources S_1 and S_2 report an event e then its denial $\bar{e}$ occurs only if both reports are false. This occurs with probability $<E_1 \rightarrow \bar{e}> \cdot <E_2 \rightarrow \bar{e}>$. The probability of the occurrance of e is given by $<E_1 \wedge E_2 \rightarrow e> = 1 - <E_1 \rightarrow \bar{e}> \cdot <E_2 \rightarrow \bar{e}>$. This can be generalized to η independent sources as follows.

$$<\bigwedge_{i=1}^{\eta} E_i \rightarrow e> = 1 - \prod_{i=1}^{\eta} <E_i \rightarrow \bar{e}>.$$

Data reliability can be incorporated into the Bayesian framework as another stage in the inference process. First, we must differentiate between the actual occurrence of a datum (D) and the report of its occurrence (D*). Assuming that the report of an event is not contingent upon which hypothesis is true, the conditional relationship between the data and the hypothesis ($<Hi \rightarrow D>$) can be decomposed into:

$$<Hi \rightarrow D^*> = <D \rightarrow D^*> <Hi \rightarrow D> + <\bar{D} \rightarrow D^*> <Hi \rightarrow \bar{D}>$$

where $<D \rightarrow D^*>$ is the probability of a report of some datum conditional upon the actual occurrence of that particular datum; $<\bar{D} \rightarrow D^*>$ is the probability of a report of some datum conditional upon the actual occurrence of any other datum; $<Hi \rightarrow \bar{D}>$ is the probability of the occurrence of any other datum conditional upon the truth of a particular hypothesis; and $<Hi \rightarrow D>$ is as defined previously. Note that $<Hi \rightarrow \bar{D}>$ equals $1 - <Hi \rightarrow D>$. Expressed in this way, the determination of the diagnostic impact of a report of some datum involves two processes, given a determination of source reliability $<D \rightarrow D^*>$: first, determination of the diagnostic impact of the reported datum $<Hi \rightarrow D>$ and the diagnostic impact of other data not reported $<Hi \rightarrow \bar{D}>$; and second, calculation of the diagnostic impact of the report $<Hi \rightarrow D^*>$ on the basis of its reliability.

2.1.1.2 Fuzzy Logic

Another valid interpretation of credibility is the degree of truth in a multi-valued, or fuzzy, logic. Fuzzy logic is developed by (Lee) from the notion of fuzzy sets of (Zadeh). A fuzzy set is a set of ordered pairs

$$A = \{(g_i/e_i)\}_1^n$$

in which g_i represents the grade (degree, or credibility) of membership of element e_i in set A, $0 \leq g_i \leq 1$, $i = 1, \ldots, n$.

Fuzzy Logic is a generalization of normal two-valued logic in which a truth value T(P) in the interval (0, 1) is assigned to each elementary proposition P in the premise of a deduction and a truth value T(C) is derived for each consequent.

The use of fuzzy logic to determine the credibility of derived results in the inference system being developed for STIS was discussed in the Final Report of the BIAS Augmentation Study (Sable).

There have recently been developed problem-solving systems based on fuzzy logic. These include FUZZY PLANNER (Kling), an extension of the problem-solving language PLANNER, and FUZZY (La Faivre) a programming language for problem-solving implemented in LISP.

2.1.1.3 Comparison of Probability Theory and Fuzzy Logic as Models for Credibility

Both the subjective probability and fuzzy logic interpretations of credibility are valid and have their legitimate roles in the overall intelligence analysis process. The credibility of atomic events (facts) can be estimated by the analyst using one or more of the following methods:

(1) directly estimated as a subjective probability,

(2) using a weighted factor tree such as that given in Figure 1-2, or

(3) using subjective assessments of a number of auxiliary probabilities such as the a-priori likelihood of the event, and the perceived reliability of the sources reporting events, compute the probability of the event in the manner of Bayes as given by (Kuhns).

The choice of which method to use should probably be left to the analyst and will depend on the confidence he feels in making the judgement and the difficulty of carrying out the implied computations. Method (2) can consider all factors, and is probably more closely allied (than is method (3)) to the method used intuitively by the analyst when he uses method (1). In method (3) certain factors will appear only implicitly in subjective probabilities rather than explicitly as in method (2). Examples are such considerations as the source's estimate of the accuracy or capability of the sensor, the signal-to-noise ratio in the observation, etc. On the other hand, the analyst may feel it is easier to make judgements of the auxiliary probabilities of method (3) than the direct judgement of credibility of method (1) or the assignment of weights of method (2). Some light may be shed on this question by experimental work in analyst performance being conducted by the Army Research Institute for the Behavioral and Social Sciences (Levine).

The above estimation of the credibilities of atomic facts and simple hypotheses is carried out quite independently of the theory of fuzzy logic. However, once these basic credibility assessments are made the situation is quite different. Fuzzy logic is the appropriate tool for placing upper bounds on the credibility of compound or derivable propositions.

Consider the structure of the Semantic Net in STIS. Here we store lists of entities which are members of various sets, either system sets or the sets defined by the values of other attributes. List intersection will be a basic operation performed to narrow the search for entities with a desired combination of properties. If we consider these lists as sets of ordered pairs (g_i/e_i) consisting of membership grade (credibility) and entity identifier, then the intersection of two lists will be a list of entities which appear on both lists, along with their membership grades. According to the rules of fuzzy logic, this will be the smaller of the two credibilities, a simple selection. For a simple example of set intersection using fuzzy logic, see Figure 2-1.

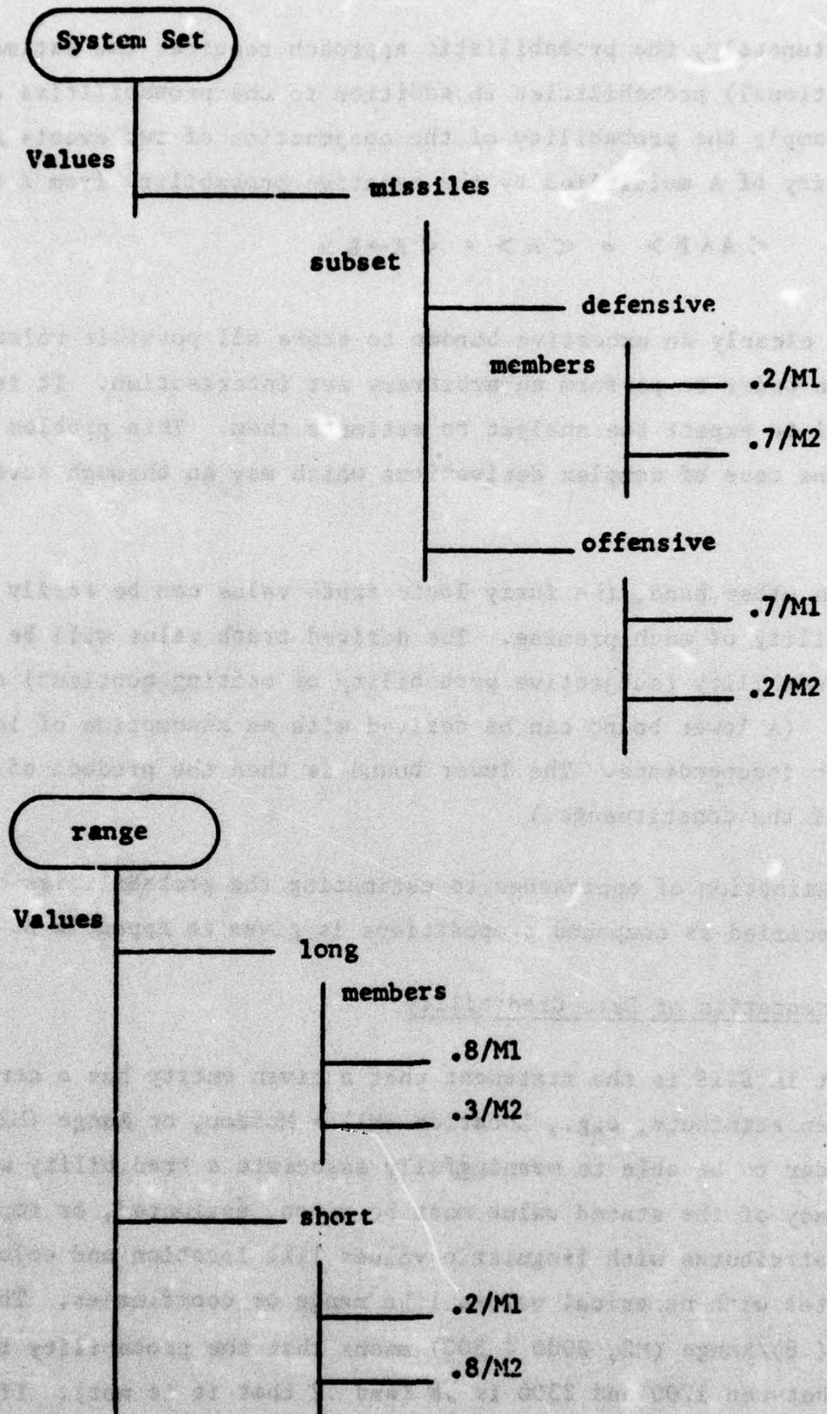

Members (long range ∧ offensive missiles) = {.7/M1, .2/M2}

Figure 2-1. Fuzzy Sets and Set Intersection

Unfortunately, the probabilistic approach requires the estimation of relative (conditional) probabilities in addition to the probabilities of atomic events. For example the probability of the conjunction of two events A and B is the probability of A multiplied by the relative probability from A to B.

$$< A \wedge B > = < A > \cdot < A \rightarrow B >$$

It is clearly an excessive burden to store all possible relative probabilities in order to perform an arbitrary set intersection. It is even more impractical to expect the analyst to estimate them. This problem is compounded in the case of complex derivations which may go through several steps.

On the other hand, the fuzzy logic truth value can be easily computed from the credibility of each premise. The derived truth value will be an upper bound of the credibility (subjective probability or betting quotient) of the derived result. (A lower bound can be derived with an assumption of logical or probabilistic independence. The lower bound is then the product of the credibilities of the constituents.)

An examination of approaches to estimating the probabilities of events which can be specified as compound propositions is given in Appendix B.

2.1.2 Representation of Data Credibility

A fact in STIS is the statement that a given entity has a certain value for a given attribute, e.g., Location (M1) = Moscow, or Range (M2) = 2000 mi. In order to be able to meaningfully associate a credibility with a fact, the accuracy of the stated value must be given, estimated, or implicit. This holds for attributes with linguistic values like location and color, as well as attributes with numerical values like range or coordinates. Thus the statement Cred (.8)/Range (M2, 2000 ± 300) means that the probability that the range of M2 is between 1700 and 2300 is .8 (and .2 that it is not). If the accuracy is not given then it may be implicit (from the missile's generic description in the STIS Entity Net) that this is a "2000 mile class" missile. This is meaningful only if the extent of the class (say 1500

to 2500 miles) is also known, and can be taken as the accuracy of the statement. Notice that if range is an indexed attribute the credibility of the set membership of M2 in a given value interval set depends on the interval as well as the fact, accuracy, and credibility. For example, if the value intervals indexed were zero to 1000, 1000 to 2000, 2000 to 3000, etc., M2 would be placed in both the 1000 to 2000, and 2000 to 3000 categories with membership grade of .5.

The specification of the value of a numerical attribute, and its tolerance, is in effect specifying the statistical distribution of values one might expect if the attribute were subject to repeated sensing. Once this distribution is established, one can then determine the probability of the value falling in a given interval. For example, if we assume that the maximum range of missile M2 is a random variable with a normal (Gaussian) error, and establish that its expected value (mean) is 2000 miles with a standard deviation of 300 miles, then the probability of a particular member of that class of missiles having a range of greater than 1700 miles is 0.84. Another way of stating this is that missile M2 is a member of the set of missiles with range greater than 1700 miles with a fuzzy set membership grade of 0.84.

This concept carries over directly to attributes which have linguistic categories rather than numbers as values. For example, the statement that the location of M1 is Moscow with credibility equal to 0.8 could be accepted to mean that M1 lies somewhere in a geographic area 80% of which is included in the geographic area known as Moscow, its grade of membership in Moscow is .8, or other consistent interpretations.

Very often, especially when the attribute is specified as a linguistic category (Moscow, red, long range, etc.), the uncertainty or credibility is also expressed as a linguistic category or probability phrase (probable, likely, certain, etc.). Recent studies (Johnson) have shown that it is feasible to convert such categories to a numerical scale with a high degree of consistency within the intelligence community, and that the encoding of these probability phrases into numerical equivalents is not appreciably influenced by sentence context.

2.1.3 Calculation of Statistical Errors of Derived or Hypothesized Values

The intelligence analyst often finds himself in the position of making a decision concerning enemy capability or technological threat based on information which is imprecisely known. (The chain from observed event to decision was illustrated in Figure 1-1.) It is as though each newly observed and reported event generates implicitly a host of questions such as those listed in Section 1.3. When the decision can be posed as a logical criterion, and where the observed or derived event parameters can be posed as statistical distributions, then it is reasonable to investigate the feasibility of associating error probabilities of Type I and Type II with each of the hypothesized decisions.

In the simplest case, consider the following scenario:

(a) a series of events has been observed (e.g., missile firings),

(b) a parameter (e.g., accuracy) is estimated as a statistical distribution,

(c) a particular interval of the value of the parameter (e.g., less than 1000 yards) is chosen as a critical region, within which a technological threat is indicated.

The question facing the analyst is whether a particular sequence of observations indicates that a technological threat exists, i.e., should he (or should he not) "sound the alarm". The possible real-world states and analyst responses are illustrated in the decision matrix shown in Table 2-1. A hypothetical distribution function for the parameter is shown in Figure 2-2. The decision as to whether or not to "sound the alarm" is a problem in decision theory and depends upon the cost of a "false alarm" or "being asleep" relative to the value of a correct decision of the situation being critical or non-critical.

TABLE 2-1. DECISION MATRIX

		ANALYST'S DECISION	
		NON-CRITICAL	CRITICAL
REAL-WORLD STATE	NON-CRITICAL	Correct	False Alarm (Type II Error)
	CRITICAL	Asleep (Type I Error)	Correct

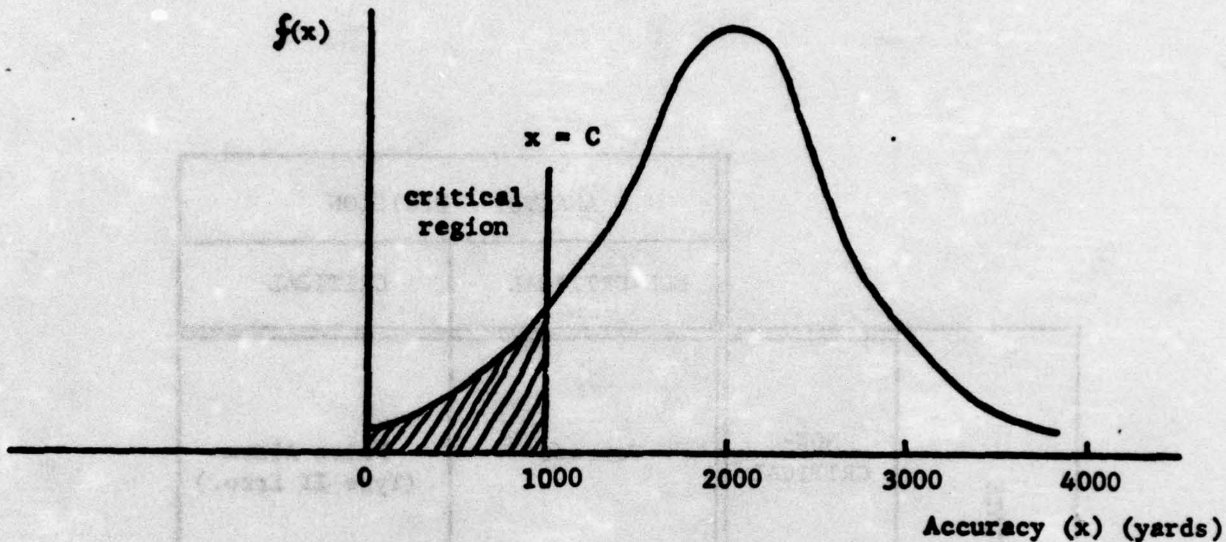

Decision	
No Alarm	Alarm
Pr{asleep} = $\int_0^C f(x)dx$	Pr{false alarm} = $\int_C^\infty f(x)dx$

Figure 2-2. Error Probabilities for Simple Decision Criterion

Unfortunately, the situation is rarely this simple. When a critical situation can be formally defined it is usually a composite condition or a functional relationship such as:

$$\text{Threat} = \left[\begin{array}{l}(\text{accuracy} < 1000 \text{ yds } \underline{\text{and}} \text{ yield} > 1 \text{ MT}) \\ \underline{\text{or}} \quad (\text{accuracy} < 2000 \text{ yds } \underline{\text{and}} \text{ yield} > 2 \text{ MT})\end{array}\right]$$

It is feasible to associate error probabilities and formal decision mechanisms for realistic criteria (such as the above) and for data values statistically derived or manually hypothesized.

2.1.4 Calculation of Credibility of Derived Information

As part of early intelligence system studies, AAI developed two algorithms for inferring (deriving) implicit information from a data base containing both facts (relations) and rules (conditional statements in a subset of the first-order predicate calculus) (Sable). The first algorithm, Subgoal Generator with Stage Preference, is a top-down problem-solver with a cost-driven search strategy, and the second, Subgoal Generator with Path Preference, is a bottom-up problem-solver with a cost-driven search strategy. Recent advances in automatic problem-solving and theorem-proving (Kuehner, Kowalski, VanderBrug) have developed ways of combining top-down and bottom-up approaches to improve search efficiency. We have taken advantage of these sdvances and have developed a detailed design of a deductive system (Goldhirsh and Carson). Along with the inference algorithms, the earlier project outlined the use of fuzzy logic, as developed by (Lee) as a potential method of determining the truth value of derived results.* This fuzzy logic truth value is an easily computed function of the credibilities of the facts and rules used in the premise of the deduction and serves as an upper bound to the credibility of the inferred result (when interpreted as a subjective probability). This avoids the computation of the subjective probability of the consequent given the probability of the antecedents, a computation that involves not only the basic credibilities but also conditional (relative) probabilities which are often unknown.

* J. Sable: Design Concept for an Augmented Relational Intelligence Analysis System. RADC-TR-73-342, pp 6-14 to 6-18 , (773189).

The use of subjective probabilities as credibilities of derived results becomes tractable if it is assumed that the atomic facts are statistically independent (an assumption which is not generally valid). In that case the computation of probability of truth depends only on the individual credibilities of the premises. This probability is in general lower than the fuzzy logic truth value and can be taken as a lower bound of the rigorous credibility of the derived result.

2.2 CONSISTENCY TESTING

The consistency of the STIS data base is a complex and, in many ways, a subtle problem which involves an interplay of each of the elements of the Concept Net: The Entity Net, the Semantic Net, and the Rules Net. The problem is engendered by the following properties of STIS and the Concept Net:

(a) Specific facts (extensional statements) are stored in the Entity Net.

(b) General rules (intentional statements) which apply to specified classes of entities and relations (or entities, attributes, and values) are stored in the Rules Net.

(c) The partial ordering (subset) relationships which exist among sets occurring as values, and the set membership relationships which exist between set terms and entities are stored in the Semantic Net.

(d) There will be a system capability to derive implicit facts from explicit rules and facts.

We say a set of statements (rules and facts) is inconsistent if we can derive from it both a new statement and its negation using valid rules of inference. This consistency condition is easy to state and in fact one can guarantee a formally consistent data base by attempting to derive the negation of each statement as it is entered into the data base. A strategy for doing this is outlined in Appendix E. However, the consistency-testing process is complicated by the set inclusion and membership relations of the Semantic Net and the temporal relationships which qualify facts in the Entity Net.

2.2.1 Range Test Operations

The analysis of whether intervals in space, time or other metric dimensions overlap is often crucial in establishing the consistency of sets of statements. Because of this the design of the node in the STIS Concept Net allows any fact (attribute/value pair in an entity) to be qualified by intervals of validity, in terms of accuracy or temporal relations.

In order to clarify these notions consider the following examples in which the relation following the slash qualifies the preceding relational statement:

(1)
- Range (M2, 2000) / Accuracy (300), Era (1972, 1974)
- Range (M2, 1700) / Accuracy (300), Era (1973, 1975)
- Range (x, y) $\wedge$ Range (x,z) $\wedge$ Overlap (Era,y,z) $\Rightarrow$ Overlap (Range, y, z)

Accuracy and Era are value qualifiers which specify the interval of validity of a value in terms of numeric range and time respectively. Overlap is a system predicate which takes three arguments, a qualifier (e.g., Accuracy or Era), and two values (or variables standing for those values). Overlap is true if the intervals of validity of the values of its second and third domain elements have a non-zero intersection with respect to the qualifier named in the first domain element. It is false if the two intervals do not overlap. Thus Example (1) is a consistent set. (It would be inconsistent if the range tolerances were ± 100.)

(2)
- (2.1) Loc (El, Philadelphia) / Era (1972, 1974)
- (2.2) Loc (El, Harrisburg) / Era (1973, 1975)
- (2.3) Loc (x,y) $\wedge$ Loc (x, z) $\wedge$ Overlap (Era, y,z) $\Rightarrow$ Overlap (Loc,y,z)

(3)
- Loc (El, Philadelphia) / Era (1972, 1974)
- Loc (El, Penna) / Era (1973, 1975)
- Loc (x,y) $\wedge$ Loc (x, z) $\wedge$ Overlap (Era, y,z) $\Rightarrow$ Overlap (Loc, y, z)

(4)
- Loc (El, Philadelphia) / Era (1972, 1973)
- Loc (El, Harrisburg) / Era (1974, 1975)
- Loc (x,y) $\wedge$ Loc (x, z) $\wedge$ Overlap (Era,y,z) $\Rightarrow$ Overlap (Loc,y,z)

Example (2) is inconsistent because by instantiating x=El, y=Philadelphia, z=Harrisburg, we can derive the negation of (2.3):

(2.4) Loc (x,y) $\wedge$ Loc (x,z) $\wedge$ Overlap (Era,y,z) $\wedge$ $\neg$ Overlap (Loc,y,z)

Example (3) is a consistent set because with the above instantiation and use of the Semantic Net we can determine that y and z (Philadelphia and Penna) have a non-zero intersection (Philadelphia is given as a subset of (Penna), and (2.4) cannot be derived.

Exsmple (4) is a consistent set because, with the same instantiation, Overlap (y,z) is false, their interval of validity (Era) does not overlap, and (2.4) cannot be derived.

The relation Overlap invokes a generalized range test operator which determines, from subset, accuracy, and temporal relationships whether there is an overlap in the interval of validity of specified statements. It is important to realize, however, that the Overlap operator must be invoked by a rule and cannot be invoked automatically since, in general, multiple values for a given attribute may not imply an anomaly. For example, defensive missiles and offensive missiles may both be subsets of the system set missiles, yet a given missile may be categorized as both offensive and defensive. That is, the subsets of a given set are not, in general, mutually exclusive (see, for example, Figure 2-1).

2.2.2. Determining Factual Anomalies with Specialized Programs

The question remains whether or not all tests for determining anomalies in factual information should invoke a consistency test using the general derivation mechanism. This may be too costly a mechanism in most instances and it may be more effective to develop specific program procedures for verifying the consistency of factual information in restricted areas of the Entity Net defined by STIS users. The determination of whether specialized fact consistency programs should be developed will depend on the efficiency of the general inference routines relative to the proposed specialized routines and on an analysis of specific consistency problems associated within a subset of the STIS data base oriented to a specific problem area. This determination will be carried out in coordination with the STIS data base managers, system developers, and users.

2.2.3 Limited Rules Consistency Testing

There is no general methodology for completely testing whether a set of rules (intentional statements) is consistent. Although the inconsistency of a set of statements can be demonstrated if it permits both a statement and its negation to be derived, there is no guaranteed general way to discover that this pair might exist. It is even impossible, in general, to conclusively state that a given statement cannot be derived, since a derivation procedure will usually be prematurely terminated due to exhausting some given resource (time and/or space) limit. (There will, of course, be instances in which it can be demonstrated conclusively, through successful exhaustion of possibilities, that a given statement cannot be derived from a given set of axioms, and is therefore deductively independent of those axioms.)

Although complete self-consistency of a set of axioms is difficult to establish in general, the restricted form in which it is proposed that rules be stated in STIS makes it possible to perform a limited type of rules consistency testing. A rule in STIS will have the form

$$\bigwedge_i A_i \Longrightarrow \bigvee_j C_j$$

That is, a conjunction of antecedents (A_i) implies a disjunction of consequents (C_j). It is expected that most rules will have a single consequent. The occurrence of relations in rules will be completely indexed in the Rule Net so that all rules which have a given relation in the antecedent or consequent can be easily retrieved and inspected. This will make it possible to detect rules which are immediately contradictory (due, for example, to inadvertent errors).

A limited rules consistency-testing program which will detect at least the following actions can be specified (in sequence):

(1) In general, each binary relation, A will have a converse B so that $B=A^{-1}$ or $A(x,y) = B(y,x)$.

One of these, say A, will be considered the primary form and A(x,y) will be substituted for B(y,x) in all rules. (Rules defining converse relations, transitivity, and set membership inheritance will be transformed and stored in special form, so that they may be invoked by special routines, more efficient than the general deductive mechanism.) This will regularize the representation of rules and reduce, by a large factor, the number which have to be stored.

(2) Using the tautologies:

$$A \Rightarrow A \vee B$$

and

$$A \wedge B \Rightarrow A$$

rules which are obviously redundant will be (at least temporarily) removed. For example, in each of the following cases the rule on the right is derivable from the rule on the left and can be removed with no loss in the deductive power of the system

$$A \Rightarrow B \qquad A \wedge C \Rightarrow B$$

$$A \Rightarrow B \qquad A \Rightarrow B \vee C$$

(That is, any expansion in the set of antecedents or consequents in a rule is redundant.)

(3) The following pair of rules are inconsistent

$$A \Rightarrow C \qquad A \Rightarrow \neg C$$

(4) For certain restricted sets of rules (to be defined as part of this task), before each rule is admitted to the set, an attempt can be made to derive the rule, and then its negation. Failure of these attempts will be indicative of the consistency and deductive independence of the rule relative to other members of the set.

(5) Finally, an additional indication of rule set consistency can be obtained at those times when new facts are successfully derived using the general inference mechanism. If, following successful derivation of a fact an attempt at deriving its negation fails, then this can be taken as a partial indication of data base consistency.

SECTION III. CREDIBILITY ANALYSIS

Data in an intelligence context reflects the analysts view of the real world state-of-affairs as derived from reports and observations. As such, each item is not known with absolute certainty but may have associated with it one or more credibility measures. These credibility factors are derived or estimated by the analyst from the reliability of the source, the accuracy of the observation, the "age" of the data, degree of independent verification, or other inputs and relationships. We view credibility as a probability of truth, or subjective probability, one of the several valid interpretations of probability, and therefore, view probability theory and inductive logic as part of the theoretical foundations of the analysis of the credibility of explicit and inferred information in an intelligence system.

3.1 MATHEMATICAL FOUNDATIONS

The problem begins with the question of the credibility of some conjecture, "fact", or hypothesis posed by the analyst or derived from a report or the data base. We use classical probability theory as the mathematical foundation for the analysis of this question. For simplicity we will call

each relational statement in the intelligence data base a fact even though it is associated with a credibility or truth probability. A fact then is a simple (atomic) relational statement in the form Rab which contains only individual constants (no uninstantiated variables). We contrast this with rules, which are universally quantified statements such as $(\forall x)$ Rax or $(\forall xyz)$ Rxy $\wedge$ Pzy $\Rightarrow$ Rxz. We will also have occasion to talk of queries or interrogations which are written as existentially quantified statements $(\exists x)$ Rex and are interpreted as "determine whether or not there are one or more instances (x_i) of x such that Rax$_i$ is a fact". As in the case of facts, rules are simply plausible statements which are associated with one or more credibility measures interpreted as measures of probability of truth or validity. Thus, we will talk of explicit facts, which are derived directly from observations and reports and exist in the data base, and implicit facts, which are derivable from explicit facts using rules and derivation procedures. Our main objective is to develop a workable model for assigning credibilities to simple and compound statements given an initial set of facts and rules (and their credibilities) and to determine under what conditions this is or is not possible.

3.1.1 Scope of the Credibility Mathematical Foundation

It is desired that our model be extendable to those facts that assign a numeric value or estimate to some attribute, such as a missile range capability. A traditional structure for information uncertainty is suggested for these facts, making use of statistical concepts, such as confidence intervals.

An adequate theory or mathematical model must be able to relate the credibility of a fact to its a-priori probability and the reliability of the observer, and observation, and report from whence that fact is derived. An adequate theory must also account for the credibility of a fact derived from other facts and general statements of stated plausibility (rules) using stated derivation procedures. One prominent reason for the selection of classical probability theory as the mathematical model for credibility is that these two transitions are best known and naturally expressed in the language of subjective (truth) probability.

3.1.2 Significance of the Selection of Classical Truth Probability for the Fact Credibility

Within the above described scope, the chief alternatives to the mathematical model selected were the use of elements of fuzzy logic instead of truth probability, and the use of a general truth value rather than a truth probability estimate. Fuzzy logic is defined so that for two statements p and q, the compound statement "p and q" is assigned a truth value which is the minimum truth value among the two separate statements p,q. In addition, the truth value is sometimes thought of as merely a helpful general indicator with no attempt to connect with the truth probability of a particular fact. We note here that we use the word "fact", in this document, without the conventional association with complete certainty.

Attractive features of these alternatives include:

(1) Simplicity of the minimum rule,

(2) There is no need for the conditional probabilities which are apt to be unavailable in practice,

(3) The difficulties of estimating probabilities are circumvented.

Viewed in this manner, the minimum rule is fully equivalent to an extreme condition among the (known or unknown) possible conditional probabilities, namely that which assigns the highest conceivable value to the probability of the compound fact "p and q". Such extreme assumptions have a strong bias towards overestimating the truth of compound facts and of any facts derived from such compound facts, using inference rules.

In the case of a probabilistic model, the credibility of the compound statement p and q would be given by the product of the credibilities of the individual statements i.e., the credibility of p times the credibility of q. This is tantamount to assuming that the statements p and q are independent, which tends to be the most reasonable assumption given the absence of any information to the contrary (see Appendix C).

The estimation of fact truth probabilities may be a significant intellectual task, but it appears to lead to a better intelligence analysis tool than an arbitrary assignment of truth values. Much of the value of a formalized or mechanized credibility and consistency model capability lies in the comparisons between variously related facts, and the truth probability estimates provide a unifying theme to make such comparisons meaningful. In addition, as already mentioned, this unifying theme enables the transition steps between the credibility knowledge of reports and of facts derived from reports, and also between the credibility knowledge of explicit system facts and of facts derived using inference rules.

3.1.3 Treatment of the Fact Credibility Foundations

In the paragraphs which follow the method of employing truth probability estimates is described and illustrated with both examples and Venn Diagrams. Probability is represented graphically by area, and area overlap furnishes the interpretation for conditonal probability, providing a concrete visualization of our credibility model.

In Appendix B we furnish a graphical interpretation of our credibility method designed to illustrate the significance of our choice of classical truth probability as compared with other choices. The work assumes the use of truth probabilities for two elementary facts, as plotted in two coordinates, the third coordinate giving probability for a compound fact based on the two elementary facts, either by simple conjunction (the "and" combination mentioned above) or other logical combinations, appearing in different figures. The result is generally a solid three dimensional plot corresponding to the scope of possible conditional probabilities between the two elementary facts.

This method illustrates, for example, how the minimum of the two truth values (from fuzzy logic), and other procedures, all are consistent with classical logic when the facts are known as certain. It further illustrates how our probability product method gives a result near the centroid of the solid truth plot, while the minimum truth value approach gives an extreme at the solid surface.

3.2 CREDIBILITY OF INDIVIDUAL FACTS AND RULES

Initially, we describe the use of probability estimates as they relate to credibility of ordinary (explicit) facts in the data base. We then expand this application in two directions. The first includes a similar credibility method for facts which are derived from ordinary facts making use of fact derivation rules. The second direction shows how the ordinary fact credibility estimates are to be obtained from the relevant reports which are the information sources, to the system.

In computing the credibility of a derived fact, contributions from both the rule, and the hypothesis facts are used.

In deriving fact credibilities from report credibility we introduce likelihood numbers, which are related to credibilities. This provides a practical way of updating fact credibility when a new relevant report enters the system. This is conditioned, of course, by the competence employed in making the report likelihood estimates.

(Here the emphasis is upon facts which fit in well with the idea of an estimated truth probability. This excludes facts which essentially assign numeric values. The method suggested for such facts is the use of an assigned best estimate with interval, to be updated using ordinary statistical methods.

3.2.1 Individual Fact Probability

Here we describe the use of ordinary probability estimates as associated with those system facts where it makes practical sense to designate the chance that a system fact is actually true. In this discussion, we make liberal use of examples and a graphical presentation of such facts, with areas representing the truth probability estimates. In Appendix B, Theoretical Foundations, a different graphical technique is employed. There the purpose is to demonstrate the connections with the procedures when the practice is to assume complete certainty in system information.

In practice, it is natural and profitable to be influenced by the whole pattern of associated facts in estimating the truth probability of any one given system fact. This larger viewpoint is discussed in Section IV, Consistency Analysis. Here we are concerned with individual non-isolated facts taken together in simple combinations.

An example of such a concern is the question of the estimate that both of two facts are simultaneously true, when there are already similar estimates for each of the two facts separately. Such problems are also in the later discussion dealing with the derivation of facts from rules.

3.2.1.1 Relationships Between Two Individual Facts

We assume that individual facts may in each case be associated with a number, ranging from 0 to 1, which represents an estimate of the probability that the fact is true. We use fact in the sense of statement, meaning that it is not necessarily true, contrary to general usage. Such statements form the bulk of the data of the system, examples of which follow:

.9, DA (1, 2)	System 1 was developed at facility 2
.8, WW (2, 7)	Person 2 works with person 7
1.0, WA (4, 3)	Person 4 works at facility 3

An interpretation of these statements can be made using traditional Venn diagrams, but with areas proportional to probability. Thus, a diagram of the first two statements may appear thus:

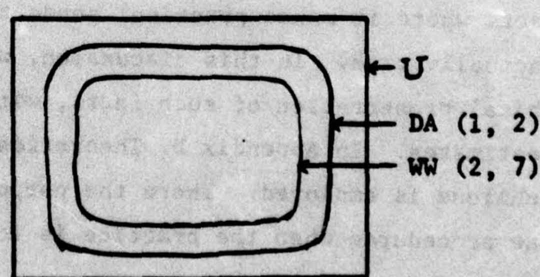

It may help to imagine a dart board in which the dart is equally liable to land at any spot in the universe U. The oval DA (1, 2) represents the fact, the dart inside DA (1, 2) represents truth, and the dart outside represents falsehood, the areas being in the proper proportions. The manner of drawing the above figure means that given WW (2, 7) is true, then surely DA (1, 2) is true. In logic notation this is WW (2, 7) $\Rightarrow$ DA (1, 2). It was not necessary, however, to represent the two facts in that manner.

Another presentation of the same two facts is as follows:

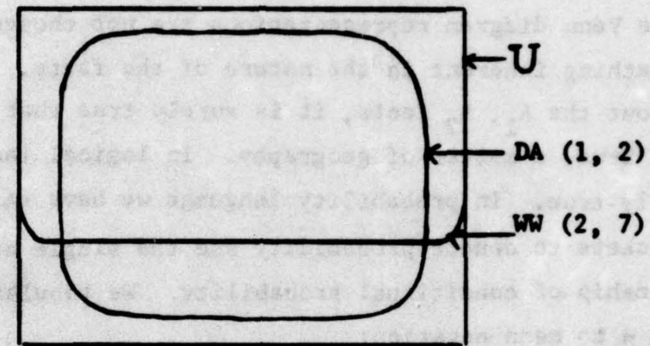

Where we have moved the WW (2, 7) oval so as to embrace the upper 80 percent of the Universe U. Under this diagram it is no longer true that WW (2, 7) $\Rightarrow$ DA (1, 2). That is to say, if the dart should land inside WW (2, 7), it may or may not land inside DA (1, 2). More specifically, there is an approximately 7 or 8 percent chance that the dart might land in that upper portion of the area corresponding to DA (1, 2) being false. We see that there is another property of our statement above and beyond the probability and the identify of the facts, namely relationships between the facts.

Further examples to illustrate relationships between facts follow in three pairs; with truth probabilities:

A_1	Scientist 3 works in Pittsburgh	.40, WA (3, 1)
A_2	Scientist 3 works in Pennsylvania	.60, WA (3, 10)
B_1	Scientist 2 works in Philadelphia	.40, WA (2, 2)
B_2	Scientist 2 works in Pittsburgh	.40, WA (2, 1)
C_1	Scientist 4 works in Albany	.40, WA (4, 5)
C_2	Scientist 4 is male	.75, ID (4, 1)

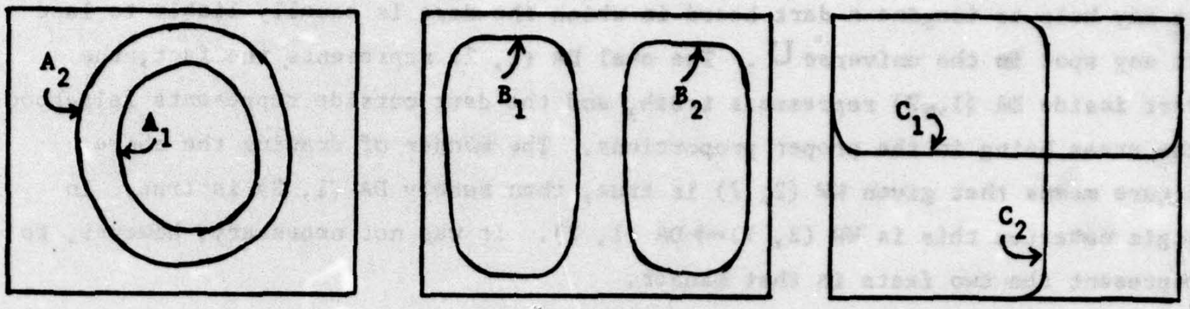

We notice that the Venn diagram representations are not chosen capriciously, but represent something inherent in the nature of the facts. Even though there is uncertainty about the A_1, A_2 facts, it is surely true that A_1 being true argues that A_2 is true, a matter of geography. In logical language we have $A_1 \Rightarrow A_2$ necessarily true. In probability language we have $\langle A_1 \rightarrow A_2 \rangle = 1$.* We use the angle brackets to denote probability and the single arrow to denote the from - to relationship of conditional probability. We tabulate the general results thus, using $\neg$ to mean negation:

	A pair	B pair	C pair
Logic Language	$A_1 \Rightarrow A_2$ $\neg A_2 \Rightarrow \neg A_1$	$B_1 \Rightarrow \neg B_2$ $B_2 \Rightarrow \neg B_1$	$\left.\begin{array}{l} C_1 \Rightarrow C_2 \\ \neg C_1 \Rightarrow C_2 \end{array}\right\}$ equally likely
Probability Language	$\langle A_1 \rightarrow A_2 \rangle = 1$ $\langle \neg A_2 \rightarrow \neg A_1 \rangle = 1$	$\langle B_1 \rightarrow B_2 \rangle = 0$ $\langle B_2 \rightarrow B_1 \rangle = 0$	$\langle C_1 \rightarrow C_2 \rangle = \langle C_2 \rangle$ $\langle C_2 \rightarrow C_1 \rangle = \langle C_1 \rangle$

It is particularly interesting to note the C pair of statements, which approximate a condition of probabilistic independence. This means, in dart language, that knowledge that the dart lands inside the C_1 oval does not lead to any expectation that the dart is inside or outside the C_2 oval, excepting such expectation as already existed. The same thought means that $C_1 \Rightarrow C_2$ and $\neg C_1 \Rightarrow C_2$ are equally likely. In the figure, the probability of failure of either implication is represented by the 12-1/2 percent portion of the total

---

* Prob$(A_1 | A_2)$ is represented as $\langle A_2 \rightarrow A_1 \rangle$ in the notation adopted here.

area at one right side corner. In more traditional probability language, the probability of the statement C_2 is not affected by the C_1 assumption: $\langle C_1 \rightarrow C_2 \rangle = \langle C_2 \rangle$. In this case we say C_2 and C_1 are independent. In popular language, we could characterize the A pair as being supportive statements, the B pair as being antagonistic, and the C pair as being impartial, or independent. Because we have in mind a system condition where conditional probabilities may be frequently unknown, the assumption of independence is the most attractive as compared with the other more extreme possibilities. Much of the substance of the following pages is an attempt to incorporate in logical inference processes such information as may be available concerning the relations between different statements, which constitute the data.

3.2.1.2 The Comprehensive Role of Conditional Probability

If we note the possibilities of simple conjunctions and disjunctions, with associated probabilities, for each of the above pairs of statements, we get the following table:

	A pair	B pair	C pair
Conjunction			
Logic	$A_1 \wedge A_2 = A_1$	$B_1 \wedge B_2$ empty	$C_1 \wedge C_2$
Probability	$\langle A_1 \wedge A_2 \rangle = \langle A_1 \rangle$	$\langle B_1 \wedge B_2 \rangle = 0$	$\langle C_1 \wedge C_2 \rangle = \langle C_1 \rangle \langle C_2 \rangle$
Disjunction			
Logic	$A_1 \vee A_2 = A_2$	$B_1 \vee B_2$	$C_1 \vee C_2$
Probability	$\langle A_1 \vee A_2 \rangle = \langle A_2 \rangle$	$\langle B_1 \vee B_2 \rangle = \langle B_1 \rangle + \langle B_2 \rangle$	$\langle C_1 \vee C_2 \rangle = 1 - (1 - \langle C_1 \rangle)(1 - \langle C_2 \rangle)$

We notice here a tendency for the assumption of probabilistic independence to be midway between the other more extreme possibilities, which we have already noted above, and which remains true for more elaborate logical expressions. A viewpoint which embraces all three conditions is the following:

$$\langle P \wedge Q \rangle = \langle P \rangle \cdot \langle P \rightarrow Q \rangle$$

The value of the conditional probability determines whether the facts P, Q are related as pair A, B, or C above, or perhaps in an intermediate manner. We tabulate the results:

$\langle Q/P \rangle$ Value	Results
1	$\langle P \wedge Q \rangle = \langle P \rangle$, as with A pair
0	$\langle P \wedge Q \rangle = 0$, as with B pair
$\langle Q \rangle$	$\langle P \wedge Q \rangle = \langle P \rangle \langle Q \rangle$, as with C pair

It is from this more general viewpoint that the remainder of this description of use of probability with data and rules is written.

3.2.1.3 <u>The Practical Advantage of the Product Rule</u>

We note, partly because of previous effort in this direction, that the work above, for pair A_1, A_2, in general, yields the results:

$$\langle A_1 \wedge A_2 \rangle = \min \{A_1, A_2\}$$
$$\langle A_2 \vee A_3 \rangle = \max \{A_1, A_2\}$$

In the above pages, for reasons of clarity, $\langle A_1 \rangle$ was selected as less than $\langle A_2 \rangle$ in agreement with the Venn diagram. But we note that these formulas are a specialization of the general conditional probability formula as discussed above, and this appears to be the source of their validity and usefulness in a probability framework. A direct application of the min-max formulas to general statements has more difficulty than serious inaccuracy where the statements do not have the properties assumed, as the A pair do. There is a failure in being well defined, even for compound statements of only modest complexity. An example follows:

$$\text{(1)} \quad \langle P \wedge \neg P \rangle = \min \{\langle P \rangle, \langle \neg P \rangle\}$$
$$= 1/2 \text{ in case } \langle P \rangle = \langle \neg P \rangle = 1/2$$

$$\text{(2)} \quad \langle P \wedge \neg P \rangle = \langle \text{empty statement} \rangle = \langle \square \rangle = 0$$

(2.4) Loc (x,y) ∧ Loc (x,z) ∧ Overlap (Era,y,z) ∧ ¬ Overlap ...
2-15

Here, P and $\neg$ P clearly have a relationship such as pair B above, and we have used the inapplicable formulas appropriate to pair A. The results show that inaccurate solutions may be derived and that the solutions depend upon the form of the compound statement. These difficulties do not appear to plague work done from the general probability viewpoint. Use of the product rule is just another special rule and can lead to these difficulties in the same way that use of the min-max rule can. However, the inaccuracies and difficulties of using the product rule (in case the conditonal probabilities are not available) tend to be less than using a special rule for a more extreme case, such as data pair A or B above. Applied to the above simple case, we illustrate ($\langle P \rangle = \langle \neg P \rangle = 1/2$ case):

min-max formula	product formula	general formula
$\langle P \wedge \neg P \rangle$	$\langle P \wedge \neg P \rangle$	$\langle P \wedge \neg P \rangle = \langle \square \rangle = 0$
$= \min\,[\langle P \rangle, \langle \neg P \rangle]$	$= \langle P \rangle \langle \neg P \rangle$	or $= \langle P \rangle \langle \neg P / P \rangle$
$= 1/2$	$= 1/4$	$= \langle P \rangle\,(0) = 0$

We hope then that we will be able to use conditional probability when it is most helpful, or at least when it is available. If it is not available, though, we plan the use of the best special rule, namely the product rule, for data that is independent in it's probabilistic nature.

3.2.2 Derivation of Facts from Rules

The simpler forms of rules which produce facts derived from the basic system facts are described. These rules have the structure of an ordinary logical implication giving a single derived fact as the result (or consequent) of an assumption based upon the truth of one or more facts in the rule hypothesis. No consideration is given to the broader class of rules which may provide semantic definition, or help implement system information structure.

The method of example and graphical illustration is used to explain a technique for including in the derived fact credibility estimate an appropriate

accounting for the intrinsic uncertainties of the rule itself, as distinct from the uncertainties of the facts in the rule hypothesis. The focal point of interest is the credibility of the derived fact itself, rather than the related concept of the truth probability of the logical implication.

This method is pursued to the point of including rules whose hypothesis facts are probabilistically independent. This is characteristic of the rule examples thus far examined, and may suffice for STIS requirements for the near future. If the rule hypothesis facts are not probabilistically independent, then the use of the traditional probability product rule is still the best procedure in case the appropriate conditional probabilities are unknown. This is a chief result of the investigation of Appendix B. and of the illustrations of paragraph 3.2.1.

Another such result is that, in the event the appropriate conditional probabilities are known, then the use of them makes a very substantial improvement in the derived fact credibility estimate. A method for so doing is explained in Appendix B.

3.2.2.1 Problem Definition for a Derived Fact Probability

An example of a rule follows:

$$WA\ (2, 4) \wedge WW\ (1, 2) \rightarrow WA\ (1, 4) \qquad \text{("instantiated" form)}$$

$$WA\ (x, y) \wedge WW\ (z, x) \Rightarrow WA\ (z, y) \qquad \text{(general form)}$$

Using language as at the start of the paper, this rule helps provide information about where person #1 works on the basis of knowledge of where co-worker person #2 works. If the statement WA (1,4) is in the data base, then the rule may not have to be used, especially if the probability estimate is high such as:

0.98, WA (1, 4)

The estimation of the probability of the result or consequent of such an implication is the central probelem still remaining here. The inference process consists of answering a given query by combinations of searching the data statements directly, and using rules to get derived data. The discussion above has described a way of treating uncertainty in the original data. Now a similar, but more involved method, is needed for derived data probability.

We suppose, first, a simple single literal hypothesis in an implication rule $p \Rightarrow q$. We regard this as completely equivalent to the statement $\neg p \vee q$. Our ultimate concern is not with the probability of the implication $\langle p \rightarrow q \rangle$, but with the probability of the consequent $\langle q \rangle$.

3.2.2.2 The Use of Conditional Probabilities for a Derived Fact

In approaching this goal, we define two numbers. The probability that the consequent is true if the antecedent is true is just the conditional probability $\langle p \rightarrow q \rangle$. The probability that the consequent is true if the antecedent is false is the conditional probability $\langle \neg p \rightarrow q \rangle$.

We therefore adopt the notation:

$$(\langle \neg p \rightarrow q \rangle, \langle p \rightarrow q \rangle) \quad p \Rightarrow q$$

An example:

$$(.2, .9) \quad p \Rightarrow q$$

signifies that $\langle p \rightarrow q \rangle = .9$
$\langle \neg p \rightarrow q \rangle = .2$

Another example follows:

$$(0, 1.0) \quad p \Rightarrow q$$

This is the case when the implication $p \Rightarrow q$ is always true (compare with data pair A discussed in paragraph 3.2.1.1).

3.2.2.3 First Estimate of the Credibility of a Derived Fact

We also note that the interpretation of the first of these rule probabilities can be given in various ways, and the ways are equivalent. Using the example (.2, .9) $P \Longrightarrow q$,

$$\langle p \Rightarrow q\rangle = .9 \quad \text{when } \langle p\rangle = 1$$
$$\langle q\rangle \quad = .9 \quad \text{when } \langle p\rangle = 1$$
$$\langle p \rightarrow q\rangle = .9 \quad \text{when } \langle p\rangle = 1$$

Here we use the single arrow $p \rightarrow q$ as the event q when p is given true. It is easy to see that, using $\langle p\rangle = 1$:

$$\langle p \rightarrow q\rangle = \langle \neg p \vee q\rangle = \langle \Box \vee q\rangle = \langle q\rangle \text{ and } \langle p \rightarrow q\rangle = \langle q\rangle.$$

Whichever of the three interpretations we regard as fundamental, it remains true that our probability number (e.g., .9 above) is an estimate of the intrinsic rule uncertainty present when the hypothesis or the input is perfectly certain. Therefore, an estimate of the consequent probability is given thus:

$$\text{estimate } \langle q\rangle = (.9)\langle p\rangle$$

We note that this can be regarded as a sort of machine operation, the rule being the machine, which cannot possibly have an output (consequent probability) of better quality than the input (antecedent probability). Thus, our .9 number is a unique estimate for the intrinsic confidence quality of the rule itself, which normally applies to a considerable number of possible facts or literals. This is illustrated by our early rule example:

$$WA\ (x, y) \wedge WW\ (z, x) \Rightarrow WA\ (z, y)$$

Where x, y, z may represent many possibilities of workers and facilities and a wide range of probabilities for the data itself. However, this $\langle q\rangle$ estimate thus far computed is not complete.

3.2.2.4 Improved Estimate of the Credibility of a Derived Fact

The missing portion of the $\langle q \rangle$ estimate thus far has to do with the possibility that the consequent may be true occasionally in spite of the hypothesis, that is when the hypothesis is false. Notice we are not considering the probability of the implication rule being true, $\langle p \Rightarrow q \rangle = \langle \neg p \vee q \rangle$, but rather the generally different probability of the consequent being true, $\langle q \rangle$. For example, a particular case of the above rule might be:

$$WA\ (2, 4) \wedge WW\ (1, 2) \Rightarrow WA\ (1, 4)$$

It might happen that the hypothesis here is false, e.g., because it is known that person 1 does not work with person 2. However, it may be true that person 1 works at facility 4, even though this is not by the power of the implication. If there were 10 facilities, then we might estimate there is a 10 percent chance that WA (1, 4) is true even though the hypothesis is definitely false. This is the part played by the .2 in the following example:

$$(.2, .9) \quad p \Rightarrow q$$

Therefore the final estimate of $\langle q \rangle$ is given by:

$$\begin{aligned} \langle q \rangle &= (.9)\langle p \rangle + (.2)(\langle \neg p \rangle) \\ &= (.9)\langle p \rangle + (.2)(1 - \langle p \rangle) \\ &= .2 + (.7)\langle p \rangle \end{aligned}$$

It may be argued that this refinement may not be worth incorporating, since it involves an estimate outside of the usual purpose of the implication rule. It may be further argued that, if the hypothesis is not satisfied, the number giving the probability of the consequent will be likely close to zero. Thus, it may seem that the difference, e.g., between:

$$.1,\ q \quad \text{and} \quad 0,\ q$$

may seem of impractical significance. But the reverse view of the matter may conceivably be misleading or dangerous. Consider for example, the result:

1.0, ¬ WA (1, 4) equivalent to 0, WA (1, 4)

which says that person 1 definitely does not work at facility 4, this answer is clearly poor if the only basis is the knowledge that the hypothesis WA (2, 4) ∧ WW (1, 2) is false.

3.2.2.5 Graphical Presentation of Derived Fact Method

A Venn diagram picture of our simple hypothesis rule:

(.2, .9) $p \Rightarrow q$

would appear as follows:

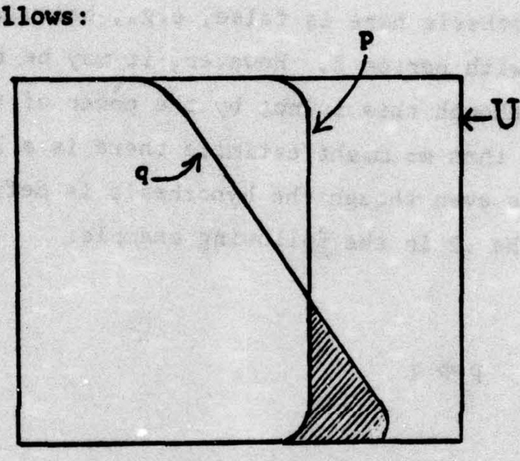

where we have shaded that portion of the <q> estimate associated with the <¬ p → q> rule probability number. If the q boundary were to coincide with the p vertical boundary, then the classical $p \Rightarrow q$ implication would be represented for the case that the implication never fails. If we use both our implication probability numbers (.9 and .2), then we have represented <q> accurately in the diagram above, given both shaded and unshaded <q> area. If we permit the refining conditional probability number to be entered as zero, then the shaded portion of the <q> area (if there is any) has been overlooked. The unshaded portion of the <q> area is that which the first probability number (.9) accounts for by itself.

3.2.2.6 Facts Derived From Compound Hypotheses

Ideas and estimates can be readily developed for implication rules with a compound hypothesis, such as:

$$p \wedge q \rightarrow r$$

If we are sure of the probability of the hypothesis, $\langle p \wedge q \rangle$, then the work of the preceeding pages may be applied by the simple device of treating $p \wedge q$ as if it were a single fact hypothesis. A much more likely situation is that we may have an estimate of the probability $\langle p \rangle$ and the probability $\langle q \rangle$ and wish to estimate $\langle r \rangle$. If nothing else is known, then $\langle p \wedge q \rangle$ may be estimated as $\langle p \rangle \langle q \rangle$, the degree of error in such an estimate depending upon the extent to which p and q really are independent. We expect the information that p and q are not probabilistically independent to appear in the system in another form such as:

$$p \rightarrow q$$

or perhaps with probability numbers as previously discussed, e.g.:

$$(.15,\ .9) \quad p \rightarrow q$$

One issue to discuss is how this relationship is to be used in estimating $\langle r \rangle$. Another issue is to see just how the $\langle r \rangle$ estimate is to be modified (as mentioned above) in the event that $p \wedge q$ is false. This last issue may be broken down into cases, such as p alone false, q alone false, etc.

3.2.3 Derivation of Facts from Reports

In many cases it may be an effective practice to estimate the credibility, or truth probability estimate, of a system fact directly. In those instances where an occasional report or a changing background condition has sufficient impact, then a fresh direct estimate may be made. However, when there is a more or less established flow of relevant reports, a more systematic approach is likely to be desirable.

Depending upon the type of the system fact, two different procedures are described for the situation of an established flow of reports. The simpler case is that where the system fact establishes some numeric estimate for a technical fact. These facts contrast with the relational type facts under discussion in recent paragraphs. It also appears practical to represent fact uncertainties with a best estimate and interval structure. The more subtle case is the application of the Bayes Theorem to update the credibility of facts recently discussed.

The numeric estimate type fact may be updated for each report by standard statistical methods. This type fact may predominate in the STIS information, but the method of updating is so common that the description has been kept brief. Depending upon particular problem features, different variants of a weighted updating procedure using a current best estimate and a current interval are anticipated.

Similar updating features for the facts with credibility involve the use of Bayes' Theorem put in the language of fact likelihood and report likelihood. These likelihoods are probability ratios which are closely linked with credibility concepts. The use of these likelihoods results in a practical directness and simplicity of method, much like the use of logarithms in certain computation problems.

The use of Bayes' Theorem in the likelihood form involves one with important operational subtleties. Questions occur involving the relative chance that a report may be received even if its content is false, and involving the extent to which a report duplicates previous reports without really offering additional support to the fact under consideration. Such problems should be faced in any event.

Accordingly, the major description of the Bayes technique is retained in the body of this report. With a continued use of examples, certain problems are further discussed in Appendix C.

3.2.3.1 Survey of the Problem

We consider first a method by which relational system facts (e.g., "Employs", "Works at", etc.) are built up from the basic report information. Previous notes have assumed facts such as:

.75, WA(2,7)	Person #2 works at facility #7 with a probability estimated at .75

and have treated the searching and combining of such facts. Here a possible procedure for obtaining and updating the .75 credibility estimate is examined. It uses the grass roots information of the reports coming into the system. In greater generality, we regard the above example as a special case of:

$$\langle e \rangle, e$$

where the event e happens to be:

$$e = WA(2,7)$$

and the credibility of the event e is:

$$\langle e \rangle = .75$$

The procedure is to consider the application of Bayes' Theorem to the reports bearing on such a fact as shown above. The central quantity which each such new report brings into the system is the likelihood or probability ratio λ, defined as follows:

$$\lambda(e \rightarrow R_e) = \frac{\langle e \rightarrow R_e \rangle}{\langle \bar{e} \rightarrow R_e \rangle}, \text{ where } R_e \text{ is the report of the event e}$$

which appears in Bayes' formulation. This is the ratio of the chances of getting the report if the fact is true, divided by the chances of getting the report if the fact is false. This estimate is made when the report is first utilized for establishing a system fact. For reports of high sharpness or discrimination, the ratio should be large.

This leads to an interesting restatement of Bayes' Theorem, making use of a similar concept of likelihood ratio for the facts themselves, designated thus:

$$L(e) = \frac{< e >}{< \bar{e} >}$$

From this viewpoint a report is exactly a likelihood ratio improver for the system facts. The restatement of Bayes Theorem is:

$$L_{new}(e) = \lambda(e \rightarrow R_e)L_{old}(e)$$

which will be proven in paragraph 3.2.3.4. Therefore, the value of the report R_e is identified with the report likelihood ratio, in the sense that the high ratio values mean the greater increase in the system fact likelihood.

Finally, a rather different sort of system fact is considered, as illustrated by:

Range (17,11$\pm$1) Missile site #17 has a range capability in the 10 to 12 mile interval.

Such facts have a different credibility structure. An essentially statistical approach is indicated for utilizing firing reports to update such a fact.

3.2.3.2 Two Reports, Bayes Theorem

Suppose we start with two reports pertinent to the fact that person #2 works at facility #7. In this simplified example, assume the reports take the form R_1 = [source S_j, fact], e.g.,

R_1 = [S_1,WA(2,7)] and

R_2 = [S_2,WA(2,7)].

Suppose further that the analyst treats R_1 alone with a credibility of .75, that is, $<R_1 \rightarrow WA(2,7)> = .75$, where the left side of the equation is the probability that the content of the R_1 report is true based upon knowledge of the R_1 report, and that report alone. We also suppose that the analyst treats R_2 alone with a credibility of .80, that is $<R_2 \rightarrow WA(2,7)> = .80$, where the left side of the equation is the probability that the R_2 report is true without any knowledge of the R_1 report. If we assume there are no further reports bearing on WA(2,7), and if we also ignore the R_2 report, the best that could be entered in the fact file of the system would be:

.75, WA(2,7)

This fact would then be a one report fact of the sort previously analyzed in a system with facts and rules. We now consider some possibilities in combining reports R_1 and R_2 to get a joint WA(2,7) result in the fact file.

A very important consideration in such a combination is the degree to which the two reports overlap. If it should happen that both R_1 and R_2 are reports, by different observers, that person #2 is listed on a facility #7 payroll listing, then the two reports are almost duplicates. The .80 credibility rating for R_2 may merely reflect the conditions of a more reliable observation. The process of combining the two reports might well approximate:

.80, WA(2,7)

thus indicating that report R_1 was essentially subsumed by R_2.

A contrasting possibility is that R_1 and R_2 may furnish essentially independent (sources S_1 and S_2 completely distinct) information supporting the fact WA(2,7). For example, R_1 may be assumed to be a payroll observation, as assumed above. But R_2 might be an observation that person #2 was seen at facility #7. Since the reports rest on a different information basis, it appears reasonable to expect a stronger supportive effect.

We attack this problem by considering Bayes' formulation of a posterior probability. This can be derived as follows:

$$\langle D \wedge H\rangle = \langle D\rangle \cdot \langle D \to H\rangle = \langle H\rangle \cdot \langle H \to D\rangle$$

$$\langle D \to H\rangle = \frac{\langle H \to D\rangle \cdot \langle H\rangle}{\langle D\rangle}$$

Let D signify data or evidence and H signify an hypothesis. Where there are only two (mutually exclusive and exhaustive) hypotheses H and $\bar{H}$

$$\langle D\rangle = \langle H \to D\rangle \cdot \langle H\rangle + \langle \bar{H} \to D\rangle \cdot \langle \bar{H}\rangle \text{ and}$$

Bayes' Theorem is written

$$\langle D\rangle \to H\rangle = \frac{\langle H \to D\rangle \cdot \langle H\rangle}{\langle H \to D\rangle \cdot \langle H\rangle + \langle \bar{H} \to D\rangle \cdot \langle \bar{H}\rangle}$$

We apply Bayes' formulation to this two report problem in the form:

$$\langle R_2 \rightarrow WA(2,7)\rangle = \frac{\langle WA(2,7) \rightarrow R_2\rangle \langle WA(2,7)\rangle}{\langle WA(2,7) \rightarrow R_2\rangle \langle WA(2,7)\rangle + \langle \overline{WA}(2,7) \rightarrow R_2\rangle \langle \overline{WA}(2,7)\rangle}$$

We also note that $\langle WA(2,7)\rangle$ is based upon R_1 above, that is to say $\langle R_1 \rightarrow WA(2,7)\rangle = \langle WA(2,7)\rangle = .75$ as commented upon above. Thus, $\langle R_2 \rightarrow WA(2,7)\rangle$ is the improved estimate of $\langle WA(2,7)\rangle$ because of the assistance of the later report R_2.

An interesting case occurs when the R_2 report is impossible when WA (2,7) is false:

$$\langle \overline{WA}(2,7) \rightarrow R_2\rangle = 0$$

This means, of course, that the R_2 report is highly reliable. Computation yields top probability for the new WA(2,7) estimate:

$$\langle R_2 \rightarrow WA(2,7)\rangle = 1$$

Thus, we see that the estimate of the likelihood ratio:

$$L\left[WA(2,7) \rightarrow R_2\right] = \frac{\langle WA(2,7) \rightarrow R_2\rangle}{\langle \overline{WA}(2,7) \rightarrow R_2\rangle}$$

is central to an understanding of the report R_2 and its effect on the $\langle WA(2,7)\rangle$ estimation. However, it is not always easy to see how this ratio is to be estimated in a given practical situation, particularly if the report foundation for the old $\langle WA(2,7)\rangle$ estimate is not known or understood. To illustrate this possibility we consider the extreme example where R_1 and R_2 are really the same report but they are accepted as distinct reports because of a clerical or technical error. It would be desirable, in this case, to have

$$\frac{\langle WA(2,7) \rightarrow R_2\rangle}{\langle \overline{WA}(2,7) \rightarrow R_2\rangle} = 1$$

because this is the value which results in the $\langle WA(2,7)\rangle$ estimate being unchanged by the introduction of the redundant R_2 report. However, the above ratio statement implies that the likelihood of getting report R_2 is completely independent of

whether WA(2,7), the hypothesis, is true or false. This sounds disturbing until one reflects that a more accurate assessment of the ratio statement is that the duplication of report R_1 is independent of the truth of the hypothesis. This degree of duplication is, of course, an important matter in an effective distilling of the power of the evidence from various reports, in support of one of the system facts. It may require a higher level of alertness of the human part of the intelligence operation than the more strictly clerical aspect of the reports.

3.2.3.3 Likelihood Ratio

We note that a second presentation of Bayes' theorem makes a reinterpretation in terms of likelihood ratio practical. We regard the quotient:

$$L(e) = \frac{\langle e \rangle}{\langle \bar{e} \rangle}$$

as definition of the betting odds on an event e or "fact" e. For example, if the fact is as likely to be false as it is to be true, then the odds for the fact is unity. On the other hand, if there is only a slight chance of the fact being false (i.e., $\langle \bar{e} \rangle \doteq 0$) then the odds of the fact is very high. This approximates the assumption of an inference system which is organized as if the data and rules are completely certain.

We also observe that the quotient:

$$\lambda(e \rightarrow R_e) = \frac{\langle e \rightarrow R_e \rangle}{\langle \bar{e} \rightarrow R_e \rangle}$$

may be regarded as the likelihood ratio of the report R_e bearing on the fact. For example, if the truth of the fact e and the falseness of the same fact are equally likely to lead to the report R_e, then the likelihood ratio of the report is unity. On the other hand, if there is very little chance of getting the report in the event that the fact is false (i.e., $\langle \bar{e} \rightarrow R_e \rangle \doteq 0$) then the report has a very high likelihood ratio. In this case, the report can be thought of as very assuring.

3.2.3.4 Restatement of Bayes' Theorem Using Likelihood Ratio

Thus, Bayes' Theorem can be restated in the following manner. The posterior (after the report) odds of a fact is equal to the prior (before the report) odds multiplied by the report likelihood ratio. It is interesting to trace the algebraic development of the viewpoint from the original, traditional presentation of the Bayes Theorem.

As presented above, Bayes' result can be stated thus:

$$\langle D \rightarrow H \rangle = \frac{\langle H \rightarrow D \rangle \cdot \langle H \rangle}{\langle D \rangle}$$

We can also write the complementary form

$$\langle D \rightarrow \bar{H} \rangle = \frac{\langle \bar{H} \rightarrow D \rangle \cdot \langle H \rangle}{\langle D \rangle}$$

Taking the ratio of the two forms, we have

$$\frac{\langle D \rightarrow H \rangle}{\langle D \rightarrow \bar{H} \rangle} = \frac{\langle H \rightarrow D \rangle}{\langle \bar{H} \rightarrow D \rangle} \cdot \frac{\langle H \rangle}{\langle \bar{H} \rangle}$$

We can call $\frac{\langle H \rangle}{\langle \bar{H} \rangle}$ the old odds on H or L_{old} (H), and $\frac{\langle D \rightarrow H \rangle}{\langle D \rightarrow \bar{H} \rangle}$ the new odds on H (considering the new datum or evidence, D) L_{new} (H). The factor $\frac{\langle H \rightarrow D \rangle}{\langle \bar{H} \rightarrow D \rangle}$ is the likelihood ratio L(H → D). With these notational shifts we have the odds/liklihood formulation of Bayes' Theorem

$$L_{new}(H) = \lambda(H \rightarrow D) \cdot L_{old}(H)$$

Using an event e as the hypothesis, and a report of event e, R_e, as the datum, we have

$$L_{new}(e) = \lambda(e \rightarrow R_e)\, L_{old}(e)$$

where we have used "L" to denote the odds of the fact, and λ the likelihood of the report R_e, as already discussed above.

Thus the likelihood ratio $\lambda(e \rightarrow R_e)$ can be viewed as a factor which transforms old odds into new odds,

where

$$\lambda(e \rightarrow R_e) = \frac{L_{new}(e)}{L_{old}(e)}$$

or two reports, R_1 and R_2, and some initial odds of an event $L_o(e)_1$ we have first

$$L_{\eta 1}(e) = \lambda_1(e \rightarrow R_1) \cdot L_o(e)$$

and then

$$L_{\eta 2}(e) = \lambda_1(e \rightarrow R_1) \cdot \lambda_2(e \rightarrow R_2) \cdot L_o(e)$$

Thus the overall likelihood ratio is the product of the individual likelihood ratios

$$L_{\eta 2} = \lambda \cdot L_o$$

where

$$\lambda = \lambda_1 \cdot \lambda_2$$

This of course can be generalized so that for η reports of an event e

$$\lambda = \prod_{i=1}^{\eta} \lambda_i(e \rightarrow R_i)$$

3.2.3.5 Value of Reports Bearing on a Given Fact

The report likelihood ratio is the essential indicator of the value of a report, and the Bayes Theorem (likelihood presentation) shows how the value of the report is realized in increasing the likelihood ratio of the fact upon which the report bears. For example, suppose we compare two possibilities:

(A) report likelihood ratio of 2,

(B) likelihood ratio of 4.

We also assume that all three reports are independent, that is, that they do not overlap significantly with each other or with the reports which have been previously assimilated into the system, resulting in a current <WA(2,7)> estimate. The reporting of (A) and of (B) are of similar value, the computation proceeding thus:

$$\text{(A)} \quad L_{new}\left[WA(2,7)\right] = (2)(2)\, L_{old}\left[WA(2,7)\right]$$

$$\text{(B)} \quad L_{new}\left[WA(2,7)\right] = (4)\, L_{old}\left[WA(2,7)\right]$$

What this means, in the event $\langle WA(2,7)\rangle_{old} = .80$, is as follows:

$$\langle WA(2,7)\rangle_{old} = .80$$

$$L_{old}\left[WA(2,7)\right] = \frac{\langle WA(2,7)\rangle_{old}}{\langle \overline{WA(2,7)}\rangle_{old}} = \frac{.80}{.20} = 4$$

$$L_{new}\left[WA(2,7)\right] = (4)(4) = 16$$

We can recover the probabilities (P) from the odds (L) as follows:

$$L = \frac{P}{\bar{P}} = \frac{P}{1-P}$$

$$P = L - LP$$

$$P\ (L+L) = L$$

$$P = \frac{L}{1+L}$$

Therefore, in the above example

$$\langle WA(2,7)\rangle_{new} = \frac{L_{new}\,[WA(2,7)]}{1 + L_{new}\,[WA(2,7)]} = \frac{16}{17} \doteq .944$$

We note that this offers a quick method of estimating the effect of a large number of reports of similar value (i.e., with similar likelihood ratios).

3.2.3.6 Operational Features, Bayes Method

As an alternative illustration, we consider again the example of two reports used before the first mention of the Bayes Theorem:

$\langle R_1 \rightarrow WA(2,7)\rangle = .75$ (No R_2 report)

$\langle R_2 \rightarrow WA(2,7)\rangle = .80$ (No R_1 report)

where the credibility number .75 is a probability estimate that the R_1 report is true, based on the R_1 report alone, and similarly for the R_2 report. We also assume that these are the only two current reports bearing on the fact WA(2,7), which is now to be brought up to date for the first time after a

long period. It is also assumed that the reports are non-overlapping (one might have been from a payroll list observation, the other from a sighting of person #2). Therefore, utilizing R_1 alone, the best that can be done is to enter in the fact file:

$$.75, WA(2,7) \quad (R_1 \text{ alone})$$

This same example was utilized in studying the importance of the report likelihood ratio:

$$L\left[WA(2,7) \rightarrow R_2\right] = \frac{\langle WA(2,7) \rightarrow R_2 \rangle}{\langle \overline{WA(2,7)} \rightarrow R_2 \rangle}$$

The subtleties in the evaluation of the ratio led to the reformulation of Bayes' theorem in terms of odds of facts and likelihood ratios of reports:

$$L_{new}(e) = \lambda(e \rightarrow R_e)\, L_{old}(e)$$

which in this case becomes:

$$\frac{\langle WA(2,7)\rangle_{new}}{\langle \overline{WA(2,7)}\rangle_{new}} = \frac{\langle WA(2,7) \rightarrow R_2\rangle}{\langle \overline{WA(2,7)} \rightarrow R_2\rangle} \; \frac{\langle WA(2,7)\rangle_{old}}{\langle \overline{WA(2,7)}\rangle_{old}}$$

where R_2 is still the second report, and the "new" and "old" estimates correspond to "after" and "before" the utilization of report R_2. The relationship is, of course, valid for any report with the appropriate understanding concerning "new" and "old". The use of likelihoods makes the problem easy, much as the introduction of logarithms in some arithmetic problems. However, it does mean that there is the need to get the problem in likelihood language. In this example that work still remains.

In the case of fact odds ratios, the translation is very easy:

$$L_{old}\left[WA(2,7)\right] = \frac{\langle WA(2,7)\rangle_{old}}{\langle \overline{WA(2,7)}\rangle_{old}}$$

$$= \frac{\langle WA(2,7)\rangle old}{1 - \langle WA(2,7)\rangle_{old}}$$

$$\langle WA(2,7)\rangle_{old} = \frac{L_{old}\left[WA(2,7)\right]}{L_{old}\left[WA(2,7)\right] + 1}$$

We see that the odds ranges from zero through all positive values, and is greater than unity for facts more likely to be true than false. In addition, any increase in the fact credibility is necessarily associated with an increase in the fact odds, and vice versa. There is no such easy transition for report likelihood ratios. Indeed, the report likelihood ratio, as presented above, is not a fraction whose denominator and numerator add up to unity, as is the case in fact odds ratios.

Now reports R_1 and R_2 have been defined in the following manner:

$$<R_1 \rightarrow WA(2,7)> = .75 \quad \text{(No use of } R_2\text{)}$$

$$<R_2 \rightarrow WA(2,7)> = .80 \quad \text{(No use of } R_1\text{)}$$

A close consideration of the matter shows that the problem is not completely defined, because there is no clear statement about the original WA(2,7) credibility estimate before either R_1 or R_2 enter the system. This follows from Bayes' formulation itself, which says that an updated fact odds (or credibility) is a result of two things, the report characteristics, and the previous fact estimates:

$$L_{new}\left[WA(2,7)\right] = \frac{<WA(2,7) \rightarrow R>}{<\overline{WA}(2,7) \rightarrow R>} \quad L_{old}\left[WA(2,7)\right]$$

The simpler (likelihood) version of the Bayes result is shown, but it is not really important whether fact credibilities or fact odds are employed. An important thing to note is that the .75 and .80 credibility estimates given above with R_1 and R_2 correspond to fact estimates on the left side of the Bayes equation, and on the right side the first factor alone represents the intrinsic report (R may be thought of as corresponding to R_1 or R_2) characteristics. The second (L_{old}) factor must be accounted for, and a clear separation of present report versus original fact estimates must be achieved.

We do this by assuming that $L_{old}\left[WA(2,7)\right]$ equals unity, tantamount to saying that the original fact was as likely true as false prior to the acceptance of either report R_1 or R_2. Other assumptions are considered in Appendix D. This means that Bayes' formulation appears thus:

$$\frac{0.75}{1-0.75} = \frac{WA(2,7) \rightarrow R_1}{\overline{WA}(2,7) \rightarrow R_1} (1) \quad \text{(No use of } R_2\text{)}$$

$$\frac{0.80}{1-0.80} = \frac{WA(2,7) \rightarrow R_2}{\overline{WA}(2,7) \rightarrow R_2} (1) \quad \text{(No use of } R_1\text{)}$$

where on the left we have translated (as already discussed) from fact credibility to fact odds. Note that there is no difficulty in considering either R_1 or R_2 as being the first accepted report. The results are that we have likelihood estimates for both reports:

$$L\left[WA(2,7) \rightarrow R_1\right] = 3 \quad \text{(No use of } R_2\text{)}$$

$$L\left[WA(2,7) \rightarrow R_2\right] = 4 \quad \text{(No use of } R_1\text{)}$$

Note that if we had assumed a higher value for the original (before either R_1 or R_2) fact odds (or credibility) then the two reports would have had lower likelihood ratios. This represents a real transfer of system information, and illustrates the significance of the above statement that the initial .75 and .80 report estimates do not completely define the problem.

The combination of the two reports is simply the routine of applying Bayes' formulation through two stages of "new" and "old". This double application of Bayes' formulation yields:

$$L_{R_1 \text{ and } R_2}\left[WA(2,7)\right] = \lambda\left[WA(2,7) \rightarrow R_2\right] \lambda\left[WA(2,7) \rightarrow R_1\right] L_{old}\left[WA(2,7)\right]$$

$$= 4 \times 3 \times 1 = 12$$

Changing from fact odds to fact credibility we obtain

$$\langle WA(2,7)\rangle_{R_1 \text{ and } R_2} = \frac{L_{R_1 \text{ and } R_2}}{L_{R_1 \text{ and } R_2} + 1}$$

$$= \frac{12}{12+1} \doteq 0.923$$

In this double Bayes' application, we have used our assumption that reports R_1 and R_2 have different sources or come from different information bases. If this is not the case, but R_2 almost duplicates R_1, then a reassessment of report likelihoods is necessary. Suppose that R_1 (with its likelihood ratio) is already assimilated in the fact estimate. Now the duplicating R_2 appears with its early report likelihood estimate of 4. This estimate of 4 (which assumes the report is not in duplication at all) is no longer acceptable, because such duplication of evidence, apparently making the associated fact appear much more credible, is quite likely without meaning. The result might then be that the system operator or analyst, from his general information perspective, decides upon an R_2 likelihood estimate (slightly above 1, perhaps) depending upon his judgment of the extent to which R_2 offers real new support to the WA(2,7) fact.

Such a situation is surely operationally important. In the case of information relating to people and technical facilities, it seems likely that many reports will occasionally have a rumor quality connecting them. That such an important reality is part of the estimation of the report and fact data system is essentially good. Of course both the estimation and the understanding of the information status strongly involve human judgment.

A record of previous reports is important, and may need review as certain new reports arrive. This is part of the more general need for a journal of all reports for general review and interpretation problems.

3.2.3.7 Direct Use of Source Veracity

When we have a direct (e.g., subjective) estimate of the credibility of a source or, more precisely, the conditional probability of an event given that we have a report of that event from a given source, then of course the credibility of that event can be taken as the credibility (veracity) of the source. For example, if we let E_1 represent a report of event e from source S_1 then the conditional probability $\langle E_1 \rightarrow e \rangle$ represents the credibility of e due to the single report. If two independent sources S_1 and S_2 report an event e then its denial $\bar{e}$ occurs only if both reports are false. This occurs

with probability $<E_1 \rightarrow \bar{e}>\cdot<E_2 \rightarrow \bar{e}>$. The probability of the occurrance of e is given by

$$<E_1 \wedge E_2 \rightarrow e> = 1-<E_1 \rightarrow \bar{e}>\cdot<E_2 \rightarrow \bar{e}>$$

This can be generalized to η independent sources as follows:

$$< \bigwedge_{i=1}^{\eta} E_i \rightarrow e> = 1-\prod_{i=1}^{\eta} <E_i \rightarrow \bar{e}>$$

3.2.3.8 Value of Reports, Decision Making

We have thus far restricted our considerations to reports which all relate to one fact in the system. In particular, the intrinsic effectiveness of the report and the degree of overlap with other reports have both entered the system with a probability estimate designated as the report likelihood ratio. We now mention some of the broader considerations affecting report value and decision making.

If two reports relate to different facts, then the relative importance of the reports will depend not only on the matters already discussed, but also on the relative importance of the affected facts, and also on their credibility estimates. That is to say, one fact may be in greater need of establishing evidence than another fact. It may well be that one fact is of considerable more importance than another. Some reports may bear on more than one di tinct fact, and may have a double or possibly treble value.

In most of these situations it appears that the nature of the investigaticn, the outlook of the investigator, and the time of the search may affect the situation so intimately that a formal structure may not be advisable for the decision process. If, for example, there is a problem in influencing the direction or quality of reports, then appropriate searches into the fact file or the record of past reports may help the analyst, but his own opinions and insights are apt to dominate the activity. Similar remarks are apt to be pertinent to the problem of the proper interpretation and disposal of a large volume of fresh reports, where some idea of the more valuable reports may be helpful.

3.2.3.9 Reports Involving Measurement Accuracy

We consider briefly the reports and facts involving measurement errors such as those resulting from observations of a missile firing range. Suppose that facts are in the system such as

Range $(17, D \pm d)$

indicating that the range of the missiles at site #17 have a minimum of D-d and a maximum of D+d. Such facts have not been explicitly considered previously because the manner in which probability is needed, and the associative data search possibilities, have both appeared simpler in nature than in the case of facts not involving measurement error, such as those concerning people and their employment.

There is, of course, an element of probability in such system facts in the degree to which D_{min} and D_{max} are assigned so as to include all possible unusual firings. This is more a matter of routine variability in observation circumstances than a matter of essential report and fact credibility. Traditional statistical approaches appear appropriate where measurement accuracy is such a factor.

One possible method is to compute a new $(D \pm d)$ pair as a result of every new report giving the Nth range reading D_N for missile site #17. A simple weighting factor α may be used in adjusting the old fact so as to include the new report R_N. For example, $\alpha = 1/10$ would mean that the new report would receive a 1/10 weight and the old fact a 9/10 weight.

Thus, a new value for the mid point of the range might be simply given thus:

$$D_{new} = (1 - \alpha)(D_{old}) + \alpha D_N$$

To get the new accuracy estimate, d_{new}, the ideas of variance and standard deviation in a normal statistical distribution may be used. Thus:

$$\text{New variance} = (1-\alpha)\left\{\left[d_{old}\right]^2 + \left[D_{new} - D_{old}\right]^2\right\} + \alpha\left[D_{new} - D_N\right]^2$$

$$d_{new} = \sqrt{\text{new variance}}$$

Many operational variations are likely in the above method. It may be that several reports may accumulate before they are incorporated in the system fact Rng ($17, D\pm d$). . There may be problems, e.g., in deciding whether site #17 is in reality two sites 17A and 17B, or not. It may be a useful practice that the ($D\pm d$) interval should be adjusted to include all but a fraction of 1% of the firings. This can be arranged by scaling "d" by a scale factor of 2 or 3 at the ends of the above routine. It may be that α should be subject to modification to reflect the conditions in different sites, or with individual reports.

It has been our purpose here to give some picture of reports and facts of a measurable engineering nature, for which missile range has served as an illustration. With suitable modification, facts concerning other properties, such as the color of smoke at a particular facility, the proportion of a certain element in some residue, may be treated in the same manner.

(The reverse of this page is blank)

SECTION IV. CONSISTENCY ANALYSIS

We examine here the question of consistency in an information system which contains propositions (explicit facts), generalized statements with quantified variables (rules), and which utilizes a system of inference to derive implicit facts, as well as retrieve explicit facts, in responding to queries.

We view the question of consistency in an intelligence system such as STIS as one of insuring that the facts, rules, and crediblities accessible to the analyst as accepted information represents a coherent set of beliefs about the real world. This is taken to mean that the system should not be able to derive deductively a proposition and its negation, and that the credibilities assigned to the propositions (facts and rules) conform to the axioms of probability theory. It can be shown (*) that unless the subjective probabilities of a set of beliefs of a given person conforms to the axiom of the probability theory, then it is possible to construct a lottery which the person always loses, independent of the true state of the world.

* F.L. Ramsey; "Truth and Probability" Pgs. 61-92 in "Kyburg"

The question of consistency will be examined in two stages. The first stage uses the system's deductive capability to attempt to assure a Concept Net of facts and rules which are deductively consistent. This phase ignores the question of credibility and is called deductive consistency. The second stage utilizes a dialog with the analyst to develop a coherent set of credibilities. This is called inductive consistency.

4.1 DEDUCTIVE CONSISTENCY

Each candidate statement (fact or rule) whether it be a query or an addition to the data base (statement corpus) can be considered as a hypothesis whose derivability or consistency relating to the data base is to be tested.

Initially, the problem can be viewed in the context of a conventional deductive logic system. Relative to some valid subset of the corpus, one of the following cases holds for any new statement.

<u>Case 1</u>. The statement is provable. Either it is

(a) explicit in the corpus, or
(b) implicit (derivable) in the corpus.

<u>Case 2</u>. The negation of the statement is provable. Either its negation is

(a) explicit in the corpus, or
(b) implicit (derivable) in the corpus.

<u>Case 3</u>. Neither the statement nor its negation is provable. In this case, the statement (or its negation) is said to be (deductively) independent of the corpus.

Thus we see the key role that the concept of a hypothesis set plays in the system. Although we have spoken of a hypothesis and its negative, more generally a mutually exclusive and exhaustive hypothesis set should be considered. A query, fact, or rule is therefore not considered in isolation, but as a member of a set of hypotheses which are mutually exclusive and exhaustive.

If the situation is that of Case 1, and the candidate statement is provable, then its admission into the corpus does not amplify the logical power

of the system. When viewed as a query it is answered in the affirmative. If it is a candidate for admission into the corpus, then the action to be taken depends on secondary objectives. If it is already in the corpus, then the occurrence of a confirming instance may be noted, possibly augmenting the credibility measure of the statement. If it is not explicit in the corpus, then the decision as to whether or not to make it explicit must be considered, basically one of a space/time tradeoff. The situation is akin to recognizing when a theorem in any deductive system is interesting, powerful or important.

If the situation is that of Case 2, and the negation of the statement is provable, then the admission of the original statement into the corpus would cause an inconsistency, leading to the probability of contradictory statements. When viewed as a query, the statement is answered in the negative. If the statement is accepted as being factual, then the corpus must be modified and rebuilt so that it is once again consistent. If the converse of the statement is explicit, then simply removing it from the corpus may be sufficient, although an attempt at derivation is required to assure that the converse cannot be derived using rules and other facts. If the converse of the accepted statement is implicit, then an examination of its derivation is required so that proper diagnosis and "surgery" can be performed.

If Case 3 holds, and neither the statement nor its negative is provable, then the statement may be considered deductively independent of the corpus and its acceptance would be a distinct amplification of its problem-solving power, deductive or otherwise. The situation is akin to adopting an axiom in a deductive system, and the questions of interest, power and importance which were raised with regard to accepting a derivable theorem, as in Case 1, are also appropriate here.

Actually, this description of Case 3 is somewhat of an oversimplification as it is, in general, theoretically impossible to determine by a mechanical procedure that a given statement in the predicate calculus is not derivable from a set of axioms. What will happen is that sometimes a statement which is, in

fact, theoretically derivable from the corpus will be said to be non-derivable because an arbitrary resource (space/time) constraint on the derivation process has been exceeded, and the process is prematurely terminated. However, in a practical sense, this situation is not catastrophic, since if the derivation is sufficiently difficult or expensive, then the incorrect answer of "no" to a question whose correct answer is "yes" has the effect of either accepting as an axiom a true statement which, in fact, is derivable (a non-injurious error), or accepting as factual a statement which is in some sense contradictory to the accepted corpus. If the situation of accepting a non-derivable statement into the corpus is always viewed as a competition among disjoint hypotheses, then there is a buffering effect due to the imposition of the (manual or machine-aided) inductive process of accepting one of a set of competing hypotheses.

On the other hand, the general non-decidability of theorems in the predicate calculus does not mean that every attempt at deriving a non-theorem will in fact exceed the resource limit. There will be many practical situations in which the derivation process will terminate in failure because the possibilities (rather than the resources) have been exhausted.

The act of choosing a member of a set of competing (consistent) hypotheses is an essential ingredient of a state-of-affairs system. It is at this point that the probabilistic nature of the corpus must be considered, and it is this aspect which allows a rational evaluation of the hypothesis set and the possible admission of a member statement into the corpus. The realm of inductive logic replaces that of deductive logic.

The situation in a state-of-affairs problem is not one in which a strict two-valued logic always applies and statements are either true or false, but one in which there is a grey-scale or continuum of truth value, or credibility, over the interval (0, 1). The credibility of some statement S will be interpreted as a probability of truth and will be written as <S>.

4.2 INDUCTIVE CONSISTENCY AS A COHERENT PATTERN OF CREDIBILITY

We view inductive consistency as a matter of comparing related information, rather than searching for logical contradictions. A coherent pattern of credibilities will be developed from a dialogue resulting from searching out facts bearing on an initial inquiry.

The illustration used here involves a pattern of credibility extended over various system facts. Starting from an assumed inquiry, these facts are related to one another by logical inference. This will then be expanded to system facts related to one another in a time or historical sense, and then to facts whose relationship is of even greater generality.

The impact of the inquiry dialogue is to illustrate the nature of consistency and the importance of effective communication between the system and the user. In later paragraphs it is emphasized that data structure design is important in enabling this communication. At that time broader patterns of system facts and inquiries, in addition to the present one, are considered.

4.2.1 Consistency Background

It appears that the use of crediblity (or subjective probability) in an inference systen changes the nature of the consistency problem. There is, of course, some variety in the nature of possible inconsistency in an inference system without the use of credibility. A direct clash in statements or derived statements may arise because of clashes in the reports made to the system, and these reported clashes may not always be ovbious. In addition, there is also the possibility of an inconsistency arising from faults in the programming or logic used in the computer, or in the use of the system made by the operator or analyst. It seems likely that the appropriate action should be to remove the observed inconsistency by improving or reinterpreting the reports made to the system, or otherwise attacking the root cause of the clash.

The introduction of credibility widens considerably the spectrum of possible inconsistencies, and brings in new subtleties. Suppose, for example, that a statement is stored in the data base with a credibility estimate

of .75, and yet the same statment is derived from other system data (facts and rules) with the estimate that the credibility of this same statement is .90. Superficially, this might seem a damaging inconsistency. On the other hand, it may be the result of using highly reliable facts with a highly credible rule to get the same statement with a much higher credibility than the estimated one when the statement was entered. Thus, not only are there new types of inconsistencies, but there are new subtleties in their interpretation and in the appropriate action, if any, to be taken. It is not inconceivable that too much consistency should be a cause for suspicion, as when witnesses to a legal case give exceedingly pat testimony. A great deal of corroboration in a state of affairs may reflect real outside circumstances but it may also reflect an organized attempt by an adversary to sell a particular misleading belief.

The essense of the inductive consistency problem can be illustrated as follows. Assume we have a set of facts which are interrelated by a set of rules of the form e $\rightarrow$ h. The premise e may be a compound proposition but for this discussion it can be assumed (without loss of generality) to be a simple statement. The situation is illustrated below.

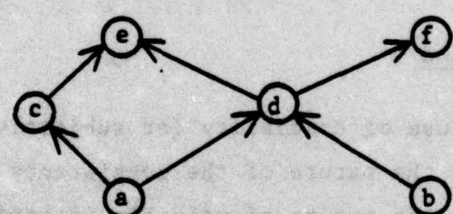

We have the facts {a,...,f} with credibilities $\{\langle a\rangle,..\langle f\rangle\}$ and the rules a $\rightarrow$ c, a $\rightarrow$ d, c $\rightarrow$ e, d $\rightarrow$ e, b $\rightarrow$ d, and d $\rightarrow$ f. Each of the rules in the form e $\rightarrow$ h has a strength given by a pair of conditional credibilities $\{\langle \bar{e} \rightarrow h\rangle, \langle e \rightarrow h\rangle\}$. The relationship among the probabilities can be derived as follows:

$$\langle h\rangle = \langle e \wedge h\rangle + \langle \bar{e} \wedge h\rangle = \langle e\rangle \langle e \rightarrow h\rangle + \langle \bar{e}\rangle \langle \bar{e} \rightarrow h\rangle$$

Thus it is clear that all the credibilities cannot be assigned independently. From the diagram it is clear that the problem is compounded by the fact that many facts can be derived through several independent paths (e.g., facts c,d,e, and f).

To illustrate this problem in its simplest form consider a report or observation E relevant to the event e. We can expand the above equation to show the new credibility of the hypothesis h, or $\langle E \rightarrow h \rangle$.

$$\langle E \rightarrow h \rangle = \langle E \rightarrow e \rangle \langle e \rightarrow h \rangle + \langle E \rightarrow \bar{e} \rangle \langle \bar{e} \rightarrow h \rangle$$

$\langle E \rightarrow e \rangle$ is the credibility of the event e in light of the new evidence E. If E is non-relevant to e then $\langle E \rightarrow e \rangle = \langle e \rangle$ and we have

$$\langle E \rightarrow h \rangle = \langle e \rangle \langle e \rightarrow h \rangle + \langle \bar{e} \rangle \langle \bar{e} \rightarrow h \rangle = \langle h \rangle$$

Thus the non-relevant evidence does not change the prior credibility of the hypothesis.

Consider the following example. Suppose we have the rule: If person a works at facility F then person a works on equipment E.

Let e = WA(a,F) (person a works at facility F.)

h = WO(a,E) (Person a works on equipment E.)

$\langle e \rangle = .2$ (There are 5 places a can work.)

$\langle h \rangle = .5$ (There are two equipments a can work on.)

Suppose the rule $e \rightarrow h$ has the strength {.1, .7}, that is $\langle \bar{e} \rightarrow h \rangle = .1$ and $\langle e \rightarrow h \rangle = .7$ (if a works at F the probability is .7 that he works on E, but if he doesn't work at F the probability is .1 that he works on E). Consider a non-relevant report E_o so that

$$\langle E_o \rightarrow e \rangle = \langle e \rangle = .2$$

Then

$$\langle E_o \rightarrow h \rangle = (.7)\ (.2) + (.1)\ (.8) = .22$$

In other words, non-relevant evidence has changed the credibility $\langle h \rangle = .5$ to $\langle E_o \rightarrow h \rangle = .22$. Now consider a relevant report E_1 with $\langle E_1 \rightarrow e \rangle = .66$. Then

$$\langle E_1 \rightarrow h \rangle = (.7)\ (.66) + (.1)\ (.33) = .495$$

Even a relevant report with credibility .66 has failed to increase the prior credibility of the hypothesis. This illustrates the fact that the credibilities $\langle e \rangle$, $\langle h \rangle$, $\langle e \rightarrow h \rangle$, and $\langle \bar{e} \rightarrow h \rangle$ cannot be assigned independently but are bound by the derived relationship. When inconsistencies in the assigned credibilities occur they should be resolved through some process.

A wide variety of consistency difficulties are considered, some of the most mathematical ones being illustrated in Appendix E. A leading technique in handling these matters is suggested as an intimate mode of dialogue between the computer and the analyst, thus making the nature of the consistency picture apparent to the analyst, and placing the interpretation and any action directly under human supervision. This is illustrated by using an example of such an analyst computer dialogue, parallel with side comments.

4.2.2 Example For Interactive Mode Of System Operation

We illustrate the more normal inconsistencies expected as a result from a query or investigation, taking, as an example, a dialogue between analyst and computer. These include clashes between ordinary STIS facts and facts derived through the inference process, both as to the nature of the fact, and as to the credibility estimate for a fact whose logical nature is unchallenged.

This dialogue example has evolved from a concept of the search process in answering a query to the system, which will be described first. The simplest situation is that in which the query can be immediately answered by searching the fact file without utilizing the inference capacity. The next simplest situation is that in which a single application of one rule is adequate to answer the query. This may also be regarded as a response using the original facts together with the first generation of derived facts. The idea may be further pursued into consideration of an additional second generation of derived facts corresponding to a second use of a rule from amongst the rules file.

The method actually contemplated for a normal query response is very different from a wholesale computing of one generation of derived facts after another. It may be that such a procedure should be used in a restricted way under appropriate circumstances, but it is apt to lead to an awkwardly bulky volume of derived statements. We anticipate using only inference rules which appear likely to generate the query answer. We expect to search only that portion of the data file containing statements of the type required by the rule in use. In the analyst/computer dialogue described below, a convenient dialogue interval is often that corresponding to the system utilization of an inference rule.

We start therefore with system information including the following facts and rules (preceded by their credibilities):

Facts

S1	0.75, WO(1,2)	person #1 works on system #2
S2	0.85, WA(1,4)	person #1 works at facility #4
S4	0.98, PO(3,2)	system #3 is part of system #2
S5	0.80, WO(6,3)	
S6	0.75, WW(7,6)	person #7 works with person #6
S9	0.70, WA(7,6)	
S10	0.95, WO(2,2)	
S11	0.95, WA(2,4)	
S12	0.65, WA(6,4)	
S21	0.90, WO(7,3)	
S22	0.95, WO(10,3)	
S23	0.90, WO(21,3)	
S24	0.95, DA(11,2)	system #11 is developed at facility #2

Rules

R4	(.05, 0.98)	$WA(x,z) \wedge WW(x,y) \Rightarrow WA(y,z)$
R6	(.02, 1.00)	$WA(x,z) \wedge WO(x,y) \Rightarrow DA(y,z)$
R7	(.05, 0.80)	$DA(y,z) \wedge PO(x,y) \Rightarrow DA(x,z)$

We also assume that the above information is illustrative of thousands of other statements in the data, and dozens of other rules. We expect, also, that many people, facilities and systems are involved in the information base. In this example we assume 300 people, 30 facilities and 45 systems.

The numbers to the left of the statements are the probability estimates of the truth of each statement. The larger number to the left of each rule is the probability estimate of the truth of the rule consequent in the case that the antecedent statements are totally certain. These have been mentioned in earlier paragraphs.

The smaller number to the left of each rule is a refinement in probability estimation mentioned in paragraph 3.2.2.2. It is the probability estimate of the truth of the rule consequent in case the antecedent of the rule is surely false. Notice, for example, that DA(7,2) can be estimated as 1/30 likely true without any antecedent at all, and without any direct data information, purely on a pot luck basis amongst 30 facilities.

4.2.3 The Investigative Dialogue

DIALOGUE	COMMENT
Q. (Analyst): What is the direct information on where system #3 is developed?	This asks for a direct search in data through facts of the form DA(3,x).
A. (Computer): None	
Q.: What is the story on first level inference information?	The rules which have conclusions matching the original query are suggested. Note one of them is an attempt to gain information through employment knowledge about people on system #3. The other rule explores through systems including system #3, and appears to have the inferior credibility but the superior search prospect. The search prospect report may be based on past search effectiveness, and on the data file completeness with respect to rule assumption facts.
A.: There are two rules: R6(.02,1.00) WA(t,x)∧WO(t,3)⇒DA(3,x). This means a search through up to 300 people at up to 30 facilities. Search prospect is fair. R7(.05,.80) DA(t,x)∧PO(3,t)⇒DA(3,x). This means a search through up to 45 systems at up to 30 facilities. Search prospect is good.	
Q.: What are the results of the searches?	Here we observe that one search uncovers a low confidence result based upon an uncertain fact about where person #6 works. The other search succeeds only in rephrasing the query in terms of a system including system #3.
A.: R6(.02,1.00) WA(6,4)∧WO(6,3)⇒DA(3,4) gives the result .53,DA(3,4) using the facts: S5 0.80, WO(6,3) S12 0.65, WA(6,4). R7(.05, .80) DA(2,x)∧PO(3,2)⇒DA(3,x) gets as far as DA(2,x), for which there is no data statement, using the statement:	

DIALOGUE

S4 0.98, PO(3,2)

Q.: Reassess the truth probability for the R6 result using the following numbers:

S5 0.85, WO(6,3)

S12 0.40, WA(6,4)

A.: .353, DA(3,4)

Q.: Ignore the R6 result. What, then, is the story on the next level of inference?

A.: The R6 search can be extended:

R6(.02,1.00) $WA(a,x) \wedge WO(a,3) \Rightarrow DA(3,x)$ using the data:

S5 0.80 WO(6,3)

S21 0.90 WO(7,3)

S22 0.95 WO(10,3)

S23 0.90 WO(21,3)

and the rule:

R4(.05,.98) $WA(t,x) \wedge WW(t,a) \Rightarrow WA(a,x)$

This means 4 searches through up to 300 people for up to 30 facilities. Search prospect is fair.

The R7 search can be extended:

R7(.05,.80) $DA(2,x) \wedge PO(3,2) \Rightarrow DA(3,x)$ using the data:

S4 0.98, PO(3,2)

and the rule:

R6(.02,1.0) $WA(t,x) \wedge WO(t,2) \Rightarrow DA(2,x)$

COMMENT

This is not a statement file change, but only a side computation. The analyst is free to explore his own conviction or interests, possibly calling upon a confidence routine not using ordinary probability ideas.

The first line of search can be continued by researching the work information on each of four persons who are stated to be working on system #3. This is possible by use of rule #4 which is based on coworker information. The search success depends upon the data file information bearing on the coworkers. Direct information bearing on one of the four persons (i.e., #6) has been of low confidence (this resulted in DA(3,4), with probability which was reassessed). The second line of search (i.e., starting with the R7 usage) can be continued by an attempt to use employment knowledge about people working on the large system #2 (includes #3). The alternate continuation of the second line is to explore systems which include #2, which in turn includes #3. The latter might be identified separately as an R7R7 search.

DIALOGUE	COMMENT
This means a search through up to 300 people for up to 30 facilities. Search prospect is fair. This R6 extension may be replaced by: R7(.05,.80) DA(t,x) $\wedge$ PO(2,t) $\Rightarrow$ DA(2,x) This means a search through up to 45 systems at up to 30 facilities. Search prospect is good.	
Q.: Ignore the R7R7 search. What, then, are the results of the second level of inference search? A.: The R6R4 search using R6(.02,1.00) WA(6,x) $\wedge$ WO(6,3) $\Rightarrow$ DA(3,x) S5 0.80, WO(6,3) R4(.05,.98) WA(7,6) $\wedge$ WW(7,6) $\Rightarrow$ WA (6,6) S6 0.75, WW(7,6) S9 0.70, WA(7,6) Gives the result <u>.442, DA(3,6)</u> The R7R6 search using: R7(.05,.80) DA(2,x) $\wedge$ PO(3,2) $\Rightarrow$ DA(3,x) S4 0.98, PO(3,2) R6(.02,1.0) WA(a,x) $\wedge$ WO(a,2) $\Rightarrow$ DA(2,x) when with data: S1 0.75, WO(1,2) S2 0.85, WA(1,4) gives the result <u>.50, DA(3,4)</u>, alternately with: S10 0.95, WO(2,2) S11 0.95, WA(2,4) gives the result <u>.70, DA(3,4)</u>	The analyst feels that information gained through subsystem of a subsystem connections are not worth pursuing for the present. The analyst now gains three results, two of which agree with the earlier result in suggesting facility #4. One of these results is the only one of the four at a modestly high (70%) probability level, and this appears to be a result of a better than average knowledge of the employment of person #2. The overall picture suggests that facility #4 is rather likely, and that the suggestion that it is facility #6 is *suspect*.

DIALOGUE	COMMENT
Q.: Enter in file: S9 .65, WA(7,6) in place of: S9 .70, WA(7,6) Investigation complete (derived data may be discarded, etc.)	After consultation with the file authority, it is agreed that the weakest link in the suspect DA(3,6) result should be placed at a lower confidence level. Data derived from an investigation of this type is normally discarded.

4.2.4 Dialogue Conclusions

This illustration concludes our investigation into consistency problems which appear likely in a normal investigatory operation. It is felt that most inconsistencies have their roots in the reports fed into the system and in the resulting system facts. It also seems likely that a chief factor in the successful handling of such problems is the effectiveness of the communication between the computing system and the operator/analyst. A better name for what we have called "consistency" might well be "a coherent pattern of crediblity:.

The consistency problems treated in Appendix E have a more special nature. Such inconsistencies are not expected to occur often in normal investigatory procedures, provided the system probability formulae and the inference rules are well designed.

4.3 INFORMATION PATTERNS

We consider here the broad patterns of information arising when time is a dominant element in the information search. Though many more facts are typically part of one of these time patterns than in the logically connected information pattern illustrated above, the searching decisions show a similarity.

Search problems and data design possibilities involving both the original reports to the system, and also facts derived from those reports, are considered. The possibilities of reorganizing system facts or even introducing

new facts from the original system reports are described, using particular examples. Credibility adjustments for aging information are also illustrated.

We begin by first giving a brief method for handling the decay of fact credibility when unsupported by new reports for protracted periods of time. Then a number of possibilities for carrying out a historical investigation are considered. The use of examples, both with employment and development type information and also with radar site type information, is dominant in the consideration of various alternatives.

These include, for example:

.99, WA(2,7)	Person #2 <u>w</u>orks <u>a</u>t facility #7 with a probability estimated at .99
Range(17, 11±1)	Missile site #17 has a range capability generally in the 10 to 12 mile interval

All of the alternatives have broad data organization implications. A complete log, recording all reports entering the system, is discussed. Periodic recording of the fact file as a protective and historical search aid is analyzed and illustrated. Examples of setting up new facts and reorganizing old information from the grass roots report level are given.

The summary suggests a practical blend of all the considered methods. The judgment in achieving this blend is of the same general character as that utilized in assigning system facts from the great bulk of auxiliary facts embedded in the system reports.

4.3.1 The Information Aging Problem

Suppose we have a well reported and virtually assured fact in the system:

.99, WA(2,7)

For days, or perhaps months, this entry may serve as a good estimate. However,

if typical employment at one facility is for approximately 4 years, and if no new reports arrive to support the WA(2,7) fact, then it becomes unrealistic to continue for many years using the same entry in the fact file. This dynamic quality in the information will depend upon the class of information: some facts will lose validity at a naturally faster rate, again supposing new reports do not appear. We consider some of the general possibilities in handling data aging problems.

It is apparent that a realistic estimate of the probability of person #2 working at facility #7 will be near zero at the end of a protracted period without new reports. This would mean certainty that person #2 is not working at facility #7, which is only likely in case the period is protracted to the time when person #2 is apt to be retired. Before retirement is likely,

.05, WA(2,7)

might be a realistic fact entry, supposing the following to be true:

(1) The original reporting bearing on the fact has aged to the point of uselessness, and no new reporting has appeared.

(2) Person #2 is very likely to be employed in some one of the 20 facilities employing people with his type of experience (facility #7 is one of them).

4.3.2 The Adjustments, Local Features

It is conceivable that individual facts will require an updating procedure which should at least include a periodic reassessment of facts which have not been reassessed because of pertinent new reports. For example, in the WA(2,7) fact illustrated above, assuming no new reports appear, a quarterly reassessment routine might result in the following time pattern:

.99, WA(2,7)	first quarter
.94, WA(2,7)	second quarter
.90, WA(2,7)	third quarter
.86, WA(2,7)	fourth quarter
.82, WA(2,7)	fifth quarter

so that at the end of about three years the WA(2,7) fact has a credibility assessment to the effect that the fact is about as likely to be false as true. The above work has assumed .05 as the zero level for the fact information, and further assumed that all credibility above the .05 depreciates quarterly at a 5% rate:

$$(.99 - .05)(.95) + .05 \doteq .94$$
$$(.94 - .05)(.95) + .05 \doteq .90$$
$$(.89 - .05)(.95) + .05 \doteq .86, \text{ etc.}$$

It may well be that other system facts will have different zero levels and different depreciation rates, but the above routine may have some applicability in accounting for the intelligence aging process.

Using the same data as already appears above, namely a zero information level at .05 credibility, and a .95 depreciation factor, a graph of the quarterly credibility is plotted below.

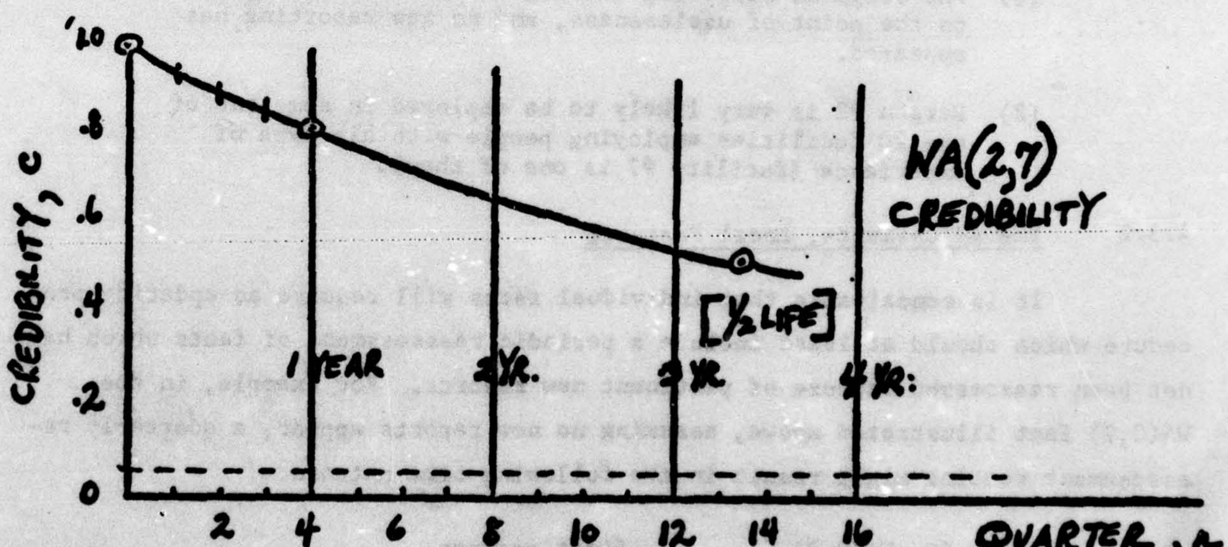

An expontial formula relating the credibility "c" to the number of quarters "n" is easy to devise. We expect the formula to follow the pattern:

$$C = .05 + .94e^{-kn} \qquad e = 2.718, \quad k = \text{constant}$$

so that we have c=.99 for η=0 and c=.05 for η very large. We notice that the .94 is the credibility above the zero information level (.94 = .99 - .05) which is subject to decay while person #2 remains employed and while no new reports are received to help identify where among the 20 facilities person #2 is employed.

The constant k can be determined by using the assumed fact that when "n" changes by 1 the decay is to a .95 proportion of the original credibility above the .05 level. Thus:

$$.95 = e^{-k}$$

$$\log_e (.95) = -k$$

$$0.0512 = k$$

and the credibility formula is:

$$c = .05 + .94e^{-0.512n}$$

This is the same sort of thing that is encountered in radioactive decay. In our case, we look for a half life for the credibility. Thus we solve for "η" when the decay factor is 1/2:

$$1/2 = e^{-kn}$$

$$\log_e (2) = kn$$

$$\frac{\log_e 2}{0.0512} \doteq n$$

$$13.5 \doteq n$$

Thus the half life is between 13 and 14 quarters, that is, between 3 and 4 years.

There are some practical motivations for exploring the above view. They concern:

(1) The basis for decay estimation,

(2) The basis for data review.

These are partly human engineering matters. It may well be that the .95 decay constant may be estimated directly, especially if such estimation is done by an experienced person. However, it may strengthen the estimation procedure to use the half life η value ($\eta \doteq 13.5$ quarters) because it has the vivid connotation of being a representative employment period at one facility.

In the same manner, the half life may furnish a very practical time increment to review the fact for possible retirement to archival storage. It is surely important to keep the on-line memory available for the most useful information. This is in keeping with the philosophy of selecting only important information from the report file for forming the system fact file.

4.3.3 Time Problems, General Features

The largely tacit assumption is often made that an investigative search concerns only the present state of affairs as understood through reports made to the system and facts deduced from these reports, together with the fact credibility estimates. The time and report problems discussed so far in this note suggest specific possibilities for keeping the system facts up to date. At times, a search into the history of a situation may be not only helpful, but possibly critically important. We consider, therefore, expanding a fact search to include temporal aspects.

Suppose then we have sufficient interest in the situation to inquire into the likelihood that person #2 was employed at facility #7, one year ago or two years ago, as well as at the present time. There are three illustrative possibilities:

(1) The employment history may be an essential part of the fact file. This means that the employment history of person #2 might be an integral part of the person #2 entity node in the data structure.

(2) A historical fact file from system use of one year ago, or also two years ago, is used. This means a general practice of periodically keeping a record of the whole fact file in a secondary memory, to facilitate such historical searches.

(3) The employment history is constructed from the record of pertinent reports received by the system. This suggests a journal-style file of reports kept in support of a ledger-style file of facts.

We intend to consider some of the advantages and disadvantages of each mode of data organization. Both because of the importance of the time element in information searches, and also because of the problems in bulk of information, the issues raised are apt to bear intimately on the system effectiveness. As a preliminary to such a consideration, we review the definitions of such words as "fact" and "report".

4.3.4 Information Definitions

In this report and in preceding technical notes the word "fact" has been used in a narrow sense, as illustrated by the following:

.99, WA(2,7)   Range (17, 11±1)

which have both appeared previously. Our facts are necessarily only primary or important facts, therefore, and this in turn is a result of human decision. Furthermore, it is in this sense that it is expected that the data bulk of the system facts is apt to be decisively less than the bulk of the system reports. This is in keeping with the previous study of the value of reports and with the patterns of analyzing the effect of various (possibly a great many) reports, all pertinent to one fact, such as the range at a particular missile site.

On the other hand, it is likely that each system report has in it what may be viewed as many auxiliary facts. The fact that missile site #17 uses missiles with a 10 to 12 mile range may be, in part, a result of one particular report which gives a time of firing and a hundred successive coordinate missile positions as coming from a radar. We never normally refer to such information as facts, even though they literally are facts. In the same spirit we keep our special usage of the word "facts", in spite of the difficulty that our facts have all levels of credibility available, whereas popular usage suggests total certainty.

A report may appear that at a particular time and place a tall man named Gerhardt, wearing a bright blue tie, got in quite a dispute with two others. This report may be one of many which support the fact:

.99, WA(2,7)

and in so doing the blue tie and the dispute may both be lost, except perhaps in a log or record of incoming system reports. It is quite conceivable that such secondary facts may remain lost, except for the possibility of a review of report information made for some special reason, such as backing up an investigative search where facility #7 is specially involved. Alternatively, it is conceivable that additional incoming reports may lead an analyst to the conclusion that the wearing of blue may have special significance, perhaps indicating an ethnic or activist group, or perhaps indicating a more specific signal. With appropriate authority he may initiate a new system fact:

.95,Tie Color (2,blue)

signifying wearing blue. In addition, he may modify selected report headers by inserting tags to make retrieval of reports bearing upon such a fact easier. Those early reports which were not so identified may require more time consuming methods in order to bring a newly created fact up to date.

When we speak of the fact file, we mean a reference organization of inportant facts together with the technique for organizing and searching through the facts, using perhaps indexing and network structures. In referring to the report log, we mean merely the time sequence of recorded reports. There may be problems in deciding whether one lengthy report should be considered as several consecutive reports. In referring to the report file, we mean the report log and also the tagging and indexing techniques which enable effective searching of the reports. This concludes the preliminary sharpening of the language before considering the three modes of data organization already mentioned, as a result of examining time problems and historical searches.

4.3.5 History Directly in the Fact File

The first mentioned possibility is to have historical information included as an essential part of the fact file, as a result of the usage made of the node network of the data structure. The degree to which this historical information might be an integral part of the fact file varies. At one extreme, it is conceivable that associated with facts such as:

.99, WA(2,7) and
Range (17, 11±1)

there might be the full record of all reports bearing on the employment of person #2 and also all reports contributing to range information of missile site #17. That is to say, such intimate report information might be repeated as part of the attribute information appearing with the entity node identifying person #2. Alternatively, the entity node identifying missile site #17 might have repeated in the attribute information the voluminous radar data which has been the source of the site range information.

The above possibility appears inefficient, partly because of the duplication of report records which will surely be necessary, in any event, as part of the report file for purposes of system protection and system reevaluation procedures, as illustrated above with the creation of the "wearing blue" fact. Another source of inefficiency is that ordinary search procedures may be impeded by the resultant bulkiness of the system fact file.

A better approach is to include only condensed or derived historical information in the fact file, using the same sort of pragmatic judgment of importance and usefulness as was illustrated in the creation of the "wearing blue" fact. One possibility, for example, is that the attribute information might have the values to make available the following facts:

.99, WA(2,7) ... Range(17, 11±1) ... (current)
.85, WA(2,7) ... Range(17, 10±1) ... (1 yr ago)
.90, WA(2,13)...(No Fact) ... (2 yrs ago)

This might mean that the entity node for person #2 will have employment location (and associated credibility) identification not only for the current time but also for the two preceeding years. Comparable comment applies to the entity node for radar site #17, where the assumption is that there is no information for the earliest year. We note that such a pattern looks good for a search into all the employees for the last two or three years, for the search appears to be a simple extension of the similar one for current employment. For a more elaborate search, such as that illustrated in paragraphs 4.2.2 and 4.2.3, special problems may arise, because inference rules are used to connect subsystem, coworkers, employment, and development location information. In such a comprehensive search it may not help much to have historical information with the WA type facts. This may not be an important difficulty because such a search is atypically difficult, and also because a modest use of historical information can well be extended to various types of facts.

A more powerful plan for history recording in the fact file is to enter the intervals of employment as attribute information associated with the entity node for person #2. Such a plan would surely be harder to implement, and also harder to utilize in the course of a historical search. Yet the search results are apt to be better, and more flexible, e.g., in answering searches concerning the state of affairs 15 months ago. We do not pursue details here, but rather continue with the second of our three general modes of data organization.

4.3.6 Use of Past Fact Files

We consider now the possibility of periodically keeping a secondary record of the whole fact file for the purpose of aiding historical searches. This is likely to mean that a more extended search involving the past state of affairs may be made, using a more complete history than is available directly in the current fact file. We suppose such records of past fact files might be available on a quarterly, or perhaps monthly basis.

There is an important distinction between the two histories that are thus considered. On the one hand, what is available might be a complete history, but one which has the viewpoint, e.g., of one year ago. In the one year since many things may have happened, including:

(1) A great volume of new reports has added to the facts and changed the old facts and their credibilities.

(2) In the case of facts which have had no reports to influence them for the last year, there is still the aging effect on the credibilities to be considered, as detailed earlier in this report.

On the other hand, the history information in the current file, if well maintained, utilizes fully all the advantages suggested in the items above, as is also true of the information about current affairs. If for example, six weeks ago a new fact has entered the fact file, either as a reassessment (as illustrated by the "wearing blue" fact) or as a change in the state of affairs, this new fact, with any appropriate history, is a part of the current fact file. Thus, the past fact file is apt to be less complete in this respect, although it is more complete with respect to the new fact's past history.

Just as keeping a time ordered log of reports entering the system has a high protective value, so too does keeping an untampered record of a past fact file. An untampered log of reports offers opportunities to check and reassess information at the grass roots level. A valid record of a past fact file protects the knowledge and insights (all too easily forgotten) which helped evolve the factual picture of that time from the report level. If it is judged that they have sufficient general value, the improvements mentioned above might be arranged in an appendix to the past fact file. The result would be that without the appendix we have the old picture of the old state of affairs. With the appendix we have the improvements to get the new picture of the old state of affairs.

We consider an example of a use of past fact files, based upon fictitious circumstances. We suppose that we have in the current fact file the following:

Range(17, 11.1 ± 0.9)

We suppose further that there is an unusual report that an odd vehicle has been sighted in the site #17 vicinity which may possibly deliver a solid propellant about every two months or so. It is also known that the rocket styles using that

propellant often have some sensitivity to the age of the fuel. To learn more of the missile site and its capability, it is decided to research the firing history and also the misfiring history. Assuming the fact files are stored away on a monthly plan, the following facts are retrieved:

Range(17, 11.1±.9) ... F(17,8,0) ... (June)
Range(17, 11.5±.9) ... F(17,14,0)... (May)
Range(17, 11.1±.9) ... F(17,15,0)... (April)
Range(17, 11.5±1.0)... F(17,14,2)... (March)

For the current month of June the above file facts mean that the bulk of the 8 firings were in the range capability interval from 10.2 miles to 12.0 miles. For the month of March the firing record shows two misfires and 14 normal firings. We also assume that each new firing has a 5% influence (as described earlier in this note) in changing the old range limits to the new range limits.

An examination of the file history lends a weak support to the possibility that the site range capability varies with a two month period. But it may well be worthwhile to reconstruct, from the system report log, the shot-by-shot range capability story. This has never heretofore been entered in the fact file. Depending upon the results of such an investigation, and upon their importance, future procedure may be to enter further facts in the fact files to improve the system performance characteristics.

In a more general sense, related investigations for other missile sites may be undertaken. The possibility of obtaining and organizing further transport information may also develop.

4.3.7 History From The Report File

In all the examples so far given the report file has appeared as a sort of last resort in the reconstruction or reinterpretation of information. Whether it is a report giving missile-firing radar position data, or a report supporting an employment or system development fact, it appears that the report file is a repository of a great number of auxiliary or secondary facts which will only rarely be referred to. The 13th x-coordinate position value in a radar sighting of a missile is not apt to be searched. Likewise a particular circumstance in the fourth reported observation supporting the fact that person #2 works at facility #7 is not likely to matter.

Yet it has seemed important to have a complete log of all reports entering the system, appropriately tagged and indexed. It may be that in some cases the report log tag identifying missile site #17 may be used to help reconstruct a more detailed firing picture than is normally kept in the fact files, because new reports have caused opinion to change about what information is important, perhaps even to the point of setting up new facts in the fact file. On the other hand, the tagging system may be largely useless, as in the odd case described above in searching for "wearing blue" information. Here the search interest may be such that nearly all reports received during a particular time interval will have to be inspected because the indexing plan did not anticipate what would be important.

4.3.8 History Implementation Summary

A method of organizing information to enable historical searches has been considered.

The foundation is a well tagged and indexed log of all reports, made as they enter the system. This has been called the report file, and may be thought of as the main journal of an information business.

The ledgers of this information business are the current fact file and the set of past fact files, possibly put in secondary memory every month; or at whatever time interval is useful. These facts constitute a much less voluminous body of data than the reports, and they are the result of human judgments of importance. It is anticipated that the bulk of investigative searching will take place over the restricted information of the fact files.

The extent to which historical information is already a part of each fact file is determined by the same sort of human judgment as is employed in initially setting up the facts of the fact file. It is expected that there will be a lot of traffic in reorganizing the fact structures and identities as new information and insights are gained. This appears to be a chief characteristic of intelligence information, and is a reason for the use of node network structures in the data organization.

(The reverse of this page is blank)

SECTION V. PROGRAM FUNCTIONAL SPECIFICATIONS

The scope of the program specifications is examined from the viewpoint of how the STIS functions are actually affected by credibility and consistency procedures. We start with charts indicating the relationships between these program specifications and preceding sections.

The result of this examination is the selection of five STIS functions to be specified because of the impact of credibility and consistency considerations. Two of these are for routine updating of the credibility of STIS facts using probability, and the routine updating of the statistical estimates for STIS numeric facts, respectively. A third function provides for a credibility aging allowance for STIS facts. A fourth function provides for the updating of an alert indicator to facilitate consistency investigations when historically related data varies in an unusual or suspicious manner. A fifth function provides for STIS investigative searches heavily conditioned by credibility considerations because of the use of fact derivation through inference rules.

All of these functions are specified at the functional level in the sense that the data states before and after the function provide the basis of description. Specific implementations and programming languages are not examined. All the functions are in some measure subject to possible adaptations and variations.

5.1 CREDIBILITY COMPUTATION (FIGURE 5-1)

We have organized the pattern of computations and decisions discussed in this report body so that a view may be had of their general relationships in the overall information plan. Fig 5-1(a) gives this picture for the raw information end of the scale; that is, from the original reports to the system facts distilled from those reports under the appropriate file authority. Fig 5-1(b) provides the picture at the more distilled end of the scale; that is, from the system facts to facts that are inferred from the original system facts.

In both Fig 5-1(a) and Fig 5-1(b), there is a division into "phase A" and "phase B", a result of considering system facts of a different nature and taking a different sort of processing. Phase A deals with system facts for which the credibility concepts of this report are applicable. For these facts it appears useful to keep track of a probability estimate (the credibility) that the fact is true. Technical employment facts are given as illustrations in the figure. Phase B deals with system facts of a numeric nature, in which it appears that uncertainty is well accounted for by a statistical method. For these facts it appears useful to keep track of a best estimate of the appropriate numeric quantity, and of its typical variability. Missile range capability or firing activity are used as illustrations in the Figure.

5.2 GENERAL INFORMATION FLOW (FIGURE 5-2)

The processors, files, and personnel for the information processing are indicated in Figure 5-2. The three figures together are a helpful guide to the program specifications which follow. The first two may also serve to illustrate preceding sectional paragraphs referenced in the Figure.

AD-A036 069 AUERBACH ASSOCIATES INC PHILADELPHIA PA F/G 9/2

THE ANALYSIS OF CREDIBILITY AND CONSISTENCY IN INTELLIGENCE DAT--ETC(U)

DEC 76 J SABLE, R DICKSON F30602-75-C-0330

UNCLASSIFIED AAI-2329-TR-1 RADC-TR-76-392 NL

2 OF 3

AD

A036069

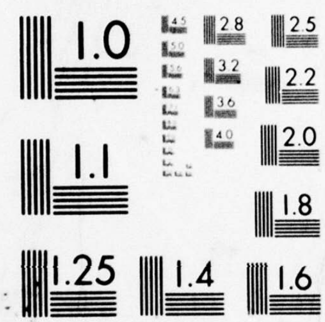

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

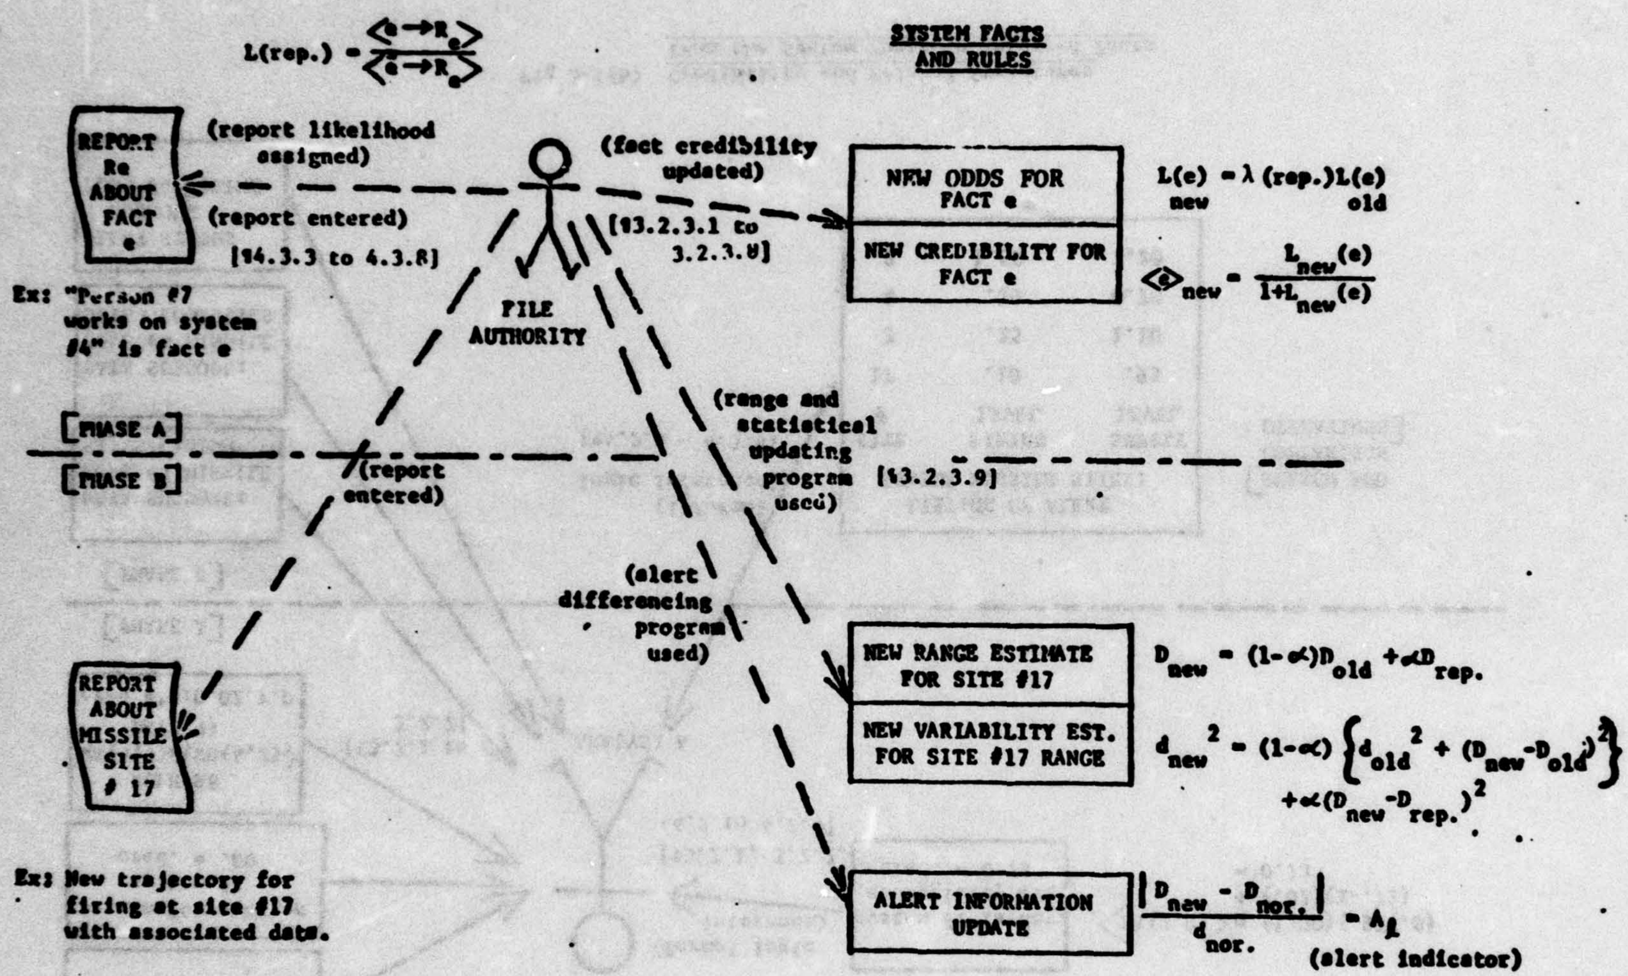

Fig 5-1(a) <u>Credibility and Related Structures From the Reports to the System Facts</u>

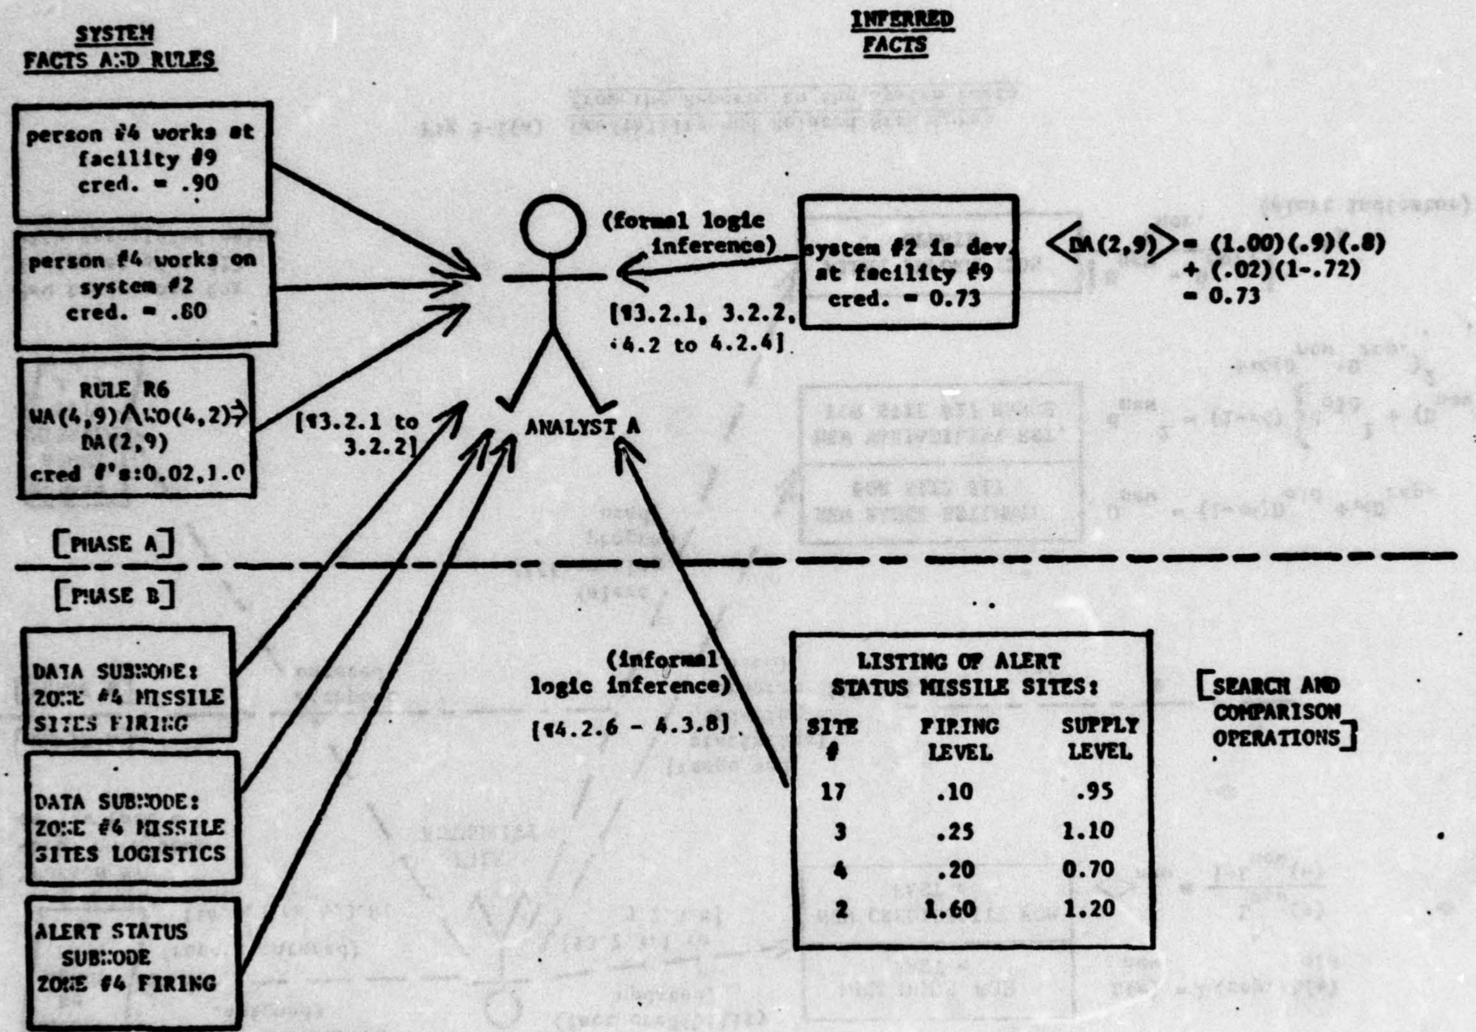

Fig 5-1(b) Credibility and Related Structures From the System Facts to Inferred Facts

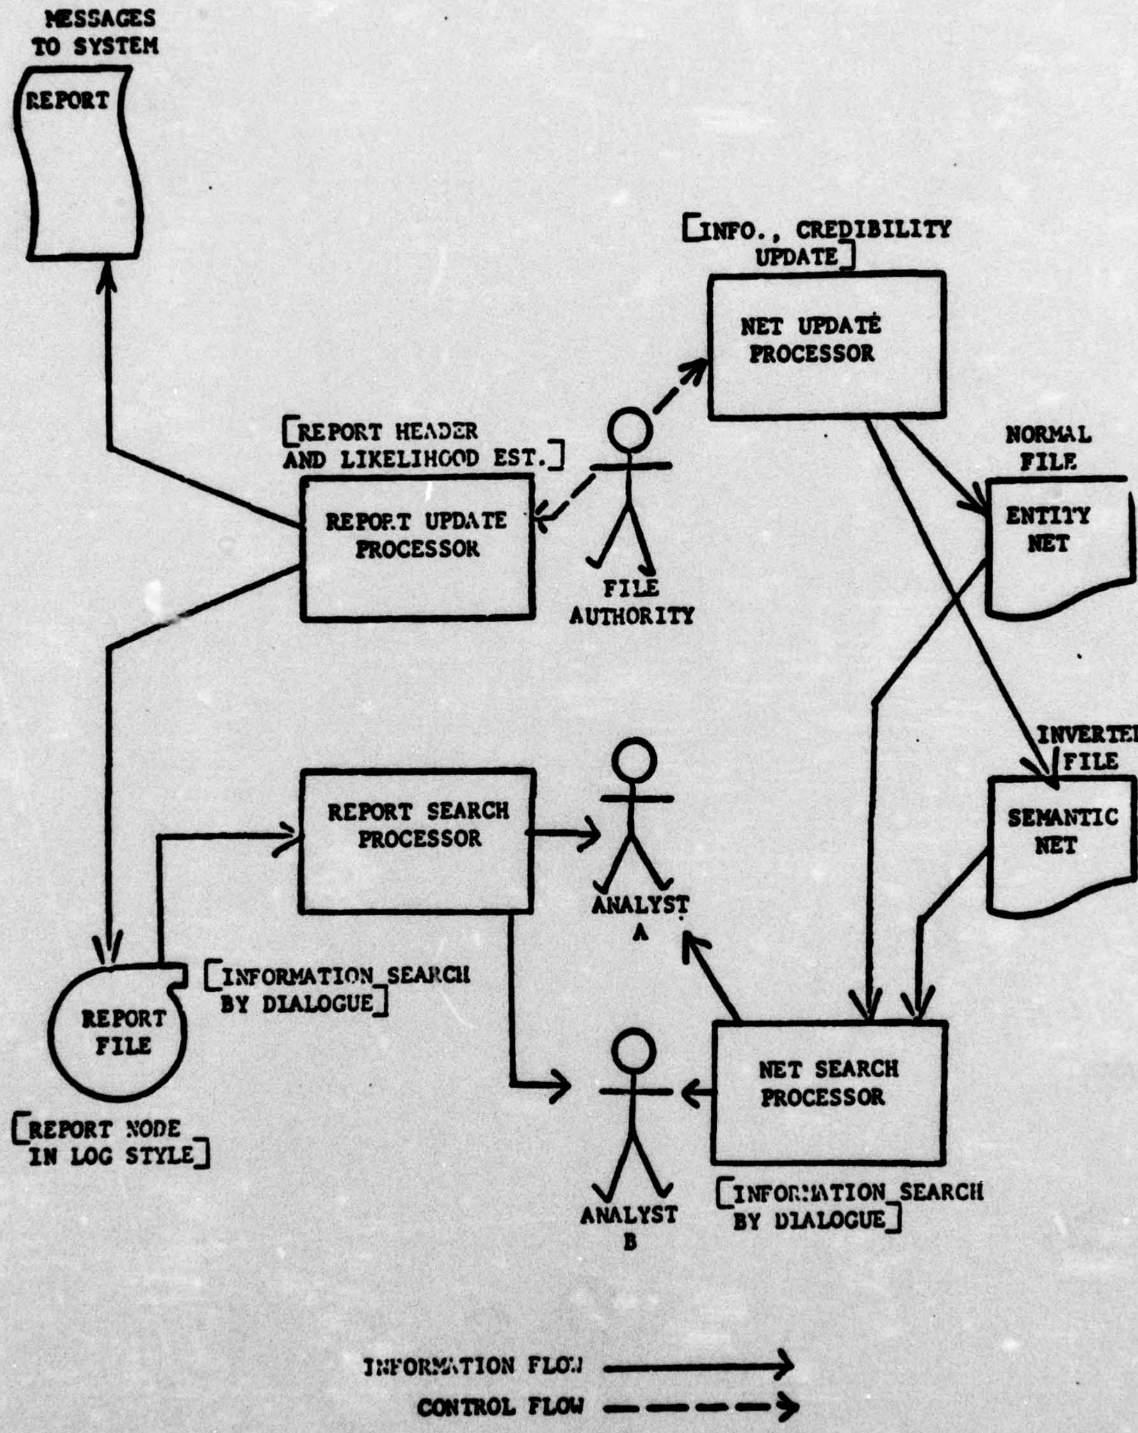

Fig 5-2 General Information Flow
(The reverse of this page is blank)

5.3 **THE PROCESSORS**

We furnish here a listing of functions to be performed by the processors appearing in Fig 5-2. It is not felt that this listing is exhaustive but consists merely of those functions mentioned in the text.

A good many of these functions appear explicitly in Fig 5-1(a) and 5-1(b). Where this does not apply, the function number appears in parentheses, and a reference is given.

It is also quite possible that some processors will have subprocessors or will consist of separate system modules. No attempt is made to resolve such possibilities, the intent being to clarify program specifications.

5.3.1 **Report Update Processor**

1. Accept likelihood estimate and report header from file authority and enter on report.

2. Accept report information and enter in report file in log manner, together with report header and likelihood estimate, and computed numeric report values.

(3) Modify selected report headers in the report file (paragraph 4.3.4).

5.3.2 **Net Update Processor**

1. Accept fact credibility or likelihood estimate and fact identification from file authority.

2. Update fact credibility, performing computation and language changes between likelihood and credibility.

(3) Update fact credibility due to aging conditions (paragraph 4.3.1, 4.3.2).

4. Compute and enter alert status update information (Note definitions below).

5. Accept numeric report values from file authority (or from report update processor) and numeric fact identification.

6. Compute and enter updated numeric fact estimates (using 5. above)

(7) Enter system information constants, such as alert threshhold values for missile activity (paragraph 4.3.1.12).

(8) Initiate and terminate system facts, with qualifying data such as credibility, time, or range of variability paragraphs 4.3.4).

5.3.3 Report Search Processor

1. Search for and give reports bearing on a particular area of interest. This may mean searching by one or more words in the report headers.

5.3.4 Net Search Processor

1. Search for and give search results bearing on a particular area of inquiry. This may mean searching selected portions of the normal and inverted files using one or more search words for search comparison purposes.

2. Search for results using logical resolution methods. This may be accomplished using a separate processing module.

5.4 DEFINITIONS AND NOTATION

We assemble here some important definitions, with associated notation, which have appeared earlier in the report body.

(1) "Credibility" of a fact means the estimate of the probability that the fact is true. This is a single number ranging from 0 to +1, which is subject to updating as appropriate reports appear. If the fact is WA(2,7) (person #2 works at facility #7), then $\langle WA(2,7) \rangle = .75$ assigns a credibility of .75, sometimes denoted thus: .75, WA(2,7).

(2) "Credibility" of a rule (i.e., a logical implication) means estimates for the probability that the conclusion of the rule is true under the assumption that the rule antecedent facts (i.e., hypothesis) are absolutely assured. These

are two numbers, each ranging from 0 to +1, one primary number giving the probability that the rule conclusion is true if the antecedent facts are true. The second, refining number, gives the probability that the rule conclusion is true if the hypothesis is false. With definition (1) notation, we suppose the rule is:

$$WA(2,7) \wedge WW(2,3) \Longrightarrow WA(3,7)$$

where we have introduced as one of the antecedent facts that person #2 works with person #3. We also assume

$$\langle WA(2,7) \rangle = .9$$

$$\langle WW(2,3) \rangle = .8$$

$$\langle WA(2,7) \wedge WW(2,3) \rangle = .9 \times .8 = .72$$

Our procedure goes thus:

$$\langle WA(3,7) \rangle = (.72)(1.0) + (1-.72)(.05)$$

$$\doteq .72 + .01 \doteq .73$$

where the two rule credibility numbers are (.05,1.0).

(3) "Odds" of a fact means the ratio of the credibility of a fact to the credibility of its negative. With definition (1) notation:

$$L[WA(2,7)] = \frac{\langle WA(2,7) \rangle}{\langle \overline{WA}(2,7) \rangle} = \frac{.75}{.25} = 3$$

(4) "Likelihood" of a report R means the ratio of the chances of getting the report if the associated fact is true, divided by the chances of getting the report if the associated fact is false. This is a single number, generally not less than unity, and is conceived of as being the essential estimate made, as the report is received, bearing on the credibility of the evidence carried by the report. When this estimate is effectively made, it naturally incorporates both the discrimination quality of the report and the degree of duplication with earlier

related reports, so that some knowledge of previous reports is inherently essential in the estimation. With definition (1) notation:

$$\lambda\left[WA(2,7) \rightarrow R\right] = \frac{\langle WA(2,7) \rightarrow R\rangle}{\langle \overline{WA(2,7)} \rightarrow R\rangle} = 2$$

meaning that under the current reporting environment, report R is twice as likely to have been received in the case that the relevant fact is true.

(5) "Numeric report values" are numbers processed from a report which present the report information in an improved or condensed manner. For example, the range capacity or the timing for a stage transition may be computed from a missile trajectory report, although neither was originally directly a part of the report. Such reports or related facts, where measure is central in the information conveyed, are referred to in phase B of the Fig 5-1 diagrams.

(6) "Alert threshhold values" are preassigned values for limits to numeric report values beyond which special alert status words enable special information explorations. For example, if the current range capability estimate or the current estimated firing activity are outside the normal operating zone, as defined by the threshhold values, then this information is posted so that the analyst is alert to the possibility of exploring information to interpret the abnormality in the operation information.

5.5 SCOPE OF PROGRAM SPECIFICATIONS

We select those functions of the processors that appear most intimately related to the area of credibility and consistency. These appear above in this manner:

A. Net Update Processor

2. Update fact credibility, performing computations and language changes between likelihood and credibility.

(3) Update fact credibility due to aging conditions (paragraphs 4.3.1, 4.3.2).

4. Compute and enter alert status update information.

6. Compute and enter updated numeric fact estimates.

B. Net Search Processor

2. Search for results using logical resolution methods. This may be accomplished using a separate processing module.

These five functions have been chosen under a broad interpretation of credibility and consistency. For example, the A.4. function above is related to the consistency of a whole group (or pattern) of facts rather than to single fact consistency problems. (Likewise, the statistical method of representing uncertainty or credibility ideas, as applied to the appropriate facts, referred to in the A.6. function, is a generalization described in paragraph 3.2.3.8.)

On the other hand, the functions listed in paragraph 5.3.1 do not appear above because the impact of credibility and consistency procedures is slight. In the case of a missile trajectory report, obtaining the range capability value (one of the numeric report values of the second function in paragraph 5.3.1) does not appear to be affected by credibility procedures. The first function in paragraph 5.3.1 may or may not have the estimate of likelihood as part of the report: otherwise this function also appears to be unaffected.

These five functions are detailed in the following paragraphs.

5.6 FUNCTIONAL PROGRAM SPECIFICATIONS FOR THE NET UPDATE PROCESSOR AND THE NET SEARCH PROCESSOR

5.6.1 Net Update Processor

5.6.1.1 Update Fact Credibility, Performing Computation and Language Changes between Likelihood and Credibility

5.6.1.1.1 Information at Start of Function

(1) $\lambda(F \rightarrow R)$ likelihood estimate is available for the evidence of report R bearing upon fact F.

(2) $L(F)_{old}$ odds estimate is available for fact F based on the evidence preceeding report R.

(3) $\langle F \rangle_{old}$ credibility estimate is available for fact F based on the evidence preceeding report R.

$$\langle F \rangle_{old} = \frac{L(F)_{old}}{1+L(F)_{old}}$$

5.6.1.1.2 Information at End of Function

(1) $\lambda(F \rightarrow R)$ likelihood estimate unchanged.

(2) $L(F)_{new}$ replaces $L(F)_{old}$ and report R is included in the fact F evidence.

(3) $\langle F \rangle_{new}$ replaces $\langle F \rangle_{old}$ and report R is included in the fact F evidence.

$$\langle F \rangle_{new} = \frac{L(F)_{new}}{1+L(F)_{new}}$$

5.6.1.1.3 Processing in the Function

(1) Compute $L(F)_{new} = L(F)_{old}\, L(F \rightarrow R)$.

(2) Compute $\langle F \rangle_{new} = \dfrac{L(F)_{new}}{1+L(F)_{new}}$.

(3) Store $L(F)_{new}$ and $\langle F \rangle_{new}$.

5.6.1.1.4 Initiation of the General Process

(1) At some point, perhaps with the first report, an $L(F)_{old}$ (and therefore $\langle F \rangle_{old}$) estimate must be made to start the process going. Care is required to make sure that no report's evidence is overlooked, or counted twice.

(2) At any point, especially when the evidence of old reports is better understood by the file authority, the file $L(F)_{old}$, $\langle F \rangle_{old}$ values may be replaced by improved values. This, in effect, reinitiates the process, and later reports may be treated routinely.

5.6.1.2 Update Fact Credibility Due to Aging Conditions (Paragraphs 4.3.1, 4.3.2)

5.6.1.2.1 Information at Start of Function

(1) D_c, system credibility depreciation constant for this type fact during a review period without reports.

(2) C_o, terminal credibility estimate for this type fact through many review periods without reports.

(3) $\langle F \rangle_{old}$ credibility estimate for fact F at start of reportless review period.

(4) $L(F)_{old}$ odds estimate for fact F at start of review period without reports.

$$L(F)_{old} = \frac{\langle F \rangle_{old}}{1-\langle F \rangle_{old}}$$

5.6.1.2.2 Information at the End of Function

(1) D_c and C_o constants unchanged

(2) $\langle F \rangle_{new}$ replaces $\langle F \rangle_{old}$

(3) $L(F)_{new}$ replaces $L(F)_{old}$

$$L(F)_{new} = \frac{\langle F \rangle_{new}}{1-\langle F \rangle_{new}}$$

5.6.1.2.3 Processing in the Function

(1) Compute $\langle F \rangle_{new} - C_o = D_c\ (\langle F \rangle_{old} - C_o)$

(2) Compute $L(F)_{new} = \dfrac{\langle F \rangle_{new}}{1-\langle F \rangle_{new}}$

(3) Store $\langle F \rangle_{new}$ and $L(F)_{new}$

5.6.1.2.4 Initiation of the General Process

(1) At any point, especially when the pattern of evidence, and gaps in reports, are better understood by the file authority, the file $\langle F \rangle_{old}$, $L(F)_{old}$ may be replaced by improved values. This, in effect, reinitiates the process, and later aging conditions may be treated routinely.

5.6.1.3 Compute and Enter Alert Status Update Information

5.6.1.3.1 Information at Start of Function

(1) $A_{\ell(old)}$ is available as alert indicator value prior to receipt of new estimate of parameter D_{new} (e.g., current range capability estimate for missile site #17)

(2) D_{new} is available as new estimate of parameter D.

(3) D_{nor} and d_{nor} are available as normal parameter value, and normal parameter tolerance estimate.

5.6.1.3.2 Information at End of Function

(1) $A_{\ell(new)}$ replaces $A_{\ell(old)}$ and the alert indicator value is based on the current estimate D_{new} of the parameter D.

(2) D_{new} is unchanged.

(3) D_{nor} and d_{nor} are unchanged.

5.6.1.3.3 Processing in the Function

(1) Compute $A_{\ell(new)} = \frac{|D_{new} - D_{nor}|}{d_{nor}}$

(2) Store $A_{\ell(new)}$

5.6.1.3.4 Initiation of the General Process

(1) As soon as system information on parameter D makes estimates for D_{nor} and d_{nor} possible and profitable, the process can be started with the current D_{new} estimate.

(2) At any point where improved knowledge warrants, improved D_{nor} and d_{nor} estimates may be employed. The process is undisturbed.

5.6.1.4 Compute and Enter Updated Numeric Fact Estimates

5.6.1.4.1 Information at Start of Function

(1) D_{rep} is available as a new reported value for the parameter D (e.g., range capability for missile site #17).

(2) D_{old} and d_{old} are available as running estimates before receipt of D_{rep}, of the value and deviation for parameter D.

(3) α is available as a smoothing or weighting constant, subject to information conditions and operator judgment.

5.6.1.4.2 Information at End of Function

(1) D_{new} and d_{new} replaces D_{old} and d_{old} as the effect of the D_{rep} value is incorporated.

(2) D_{rep} and α are unchanged.

5.6.1.4.3 Processing in the Function

(1) Compute $D_{new} = (1-\alpha)\, D_{old} + \alpha D_{rep}$

(2) Compute:

$$d_{new} = \sqrt{(1-\alpha)\,[(d_{old})^2 + (D_{new} - D_{old})^2] + \alpha[D_{new} - D_{rep}]^2}$$

(3) Store D_{new} and d_{new}

5.6.1.4.4 Initiation of the General Process

(1) A preliminary pair of estimates, before receipt of D_{rep} values, or shortly after, may be employed for initial D_{old}, d_{old} values.

(2) Increased experience may readily lead to modification of α parameter.

5.6.2 Net Search Processor

5.6.2.1 Search for Results Using Logical Resolution Methods

This may be accomplished using a separate processing module.

5.6.2.1.1 Information at Start and at End of Function

(1) <p>, p is available as fact p with current credibility estimate <p>.

(2) <q>, q is available as fact q with current credibility estimate <q>.

(3) (<r->, <r+>) $p \wedge q \Rightarrow r$ is available as rule yielding derived fact r. <r-> and <r+> are the credibility estimates for fact r in case the rule hypothesis is false or true, respectively.

5.6.2.1.2 Processing in the Function

(1) Compute <r> = <r-> (1-<p> <q>) + <r+> (<p> <q>)

(2) Prepare <r>, r for display to operator

5.6.2.1.3 Display of Search Results

(1) Search result <r>, r where r is a derived fact and <r> its credibility estimate.

(2) Background data of paragraph 5.6.2.1.1

5.6.2.1.3 Initiation of the General Process

(1) Both facts p and q are apt to be obtained as the result of a search and comparison operation in which credibility and consistency structures are insignificant. An illustration appears in paragraphs 4.2.2 and 4.2.3.

(2) Such rules may be applied consecutively, under analyst control, as illustrated in above mentioned paragraphs.

(The reverse of this page is blank)

6.1 LISP MODIFICATIONS

The LISP language suggested itself as a convenient tool for the specification, design, and implementation of information structure manipulation techniques associated with intelligence information. It would allow for the rapid development of inference strategies and credibility analysis techniques based on different information structures.

Although the LISP language structure is sufficient, certain deficiencies existed in the LISP processor available for the UNIVAC 1110 system at FTD. These deficiencies would limit the accommodation of the expected range of data access requirements for the intelligence system at FTD. The main areas requiring investigation can be summarized as follows:

(1) The transparent transfer of "pages" of in-process list structures between core and mass storage in order to increase the amount of list structure that can be accommodated in a given working storage area,

(2) The transfer and restructuring of complete information modules between core and permanent mass storage in order to efficiently utilize blocks of mass storage and provide for an on-going data base common to a group of users,

(3) Communication between LISP and non-LISP program modules,

(4) The limit on the amount of information which can be referenced in the current 1100 LISP address space, and

(5) The implementation of double precision real numbers.

The investigation and subsequent modifications to the 1100 LISP system were oriented toward producing a new tool to be utilized in a complex and high-volume data base environment. The modifications preserve the LISP language processing integrity wherever possible and are considered to be generally useful outside of the STIS environment.

6.2.1 Software Paging

In order to reduce the real memory requirements placed on the FTD 1110 computer by a LISP system, the concept of a paging environment was examined and implemented.

The most pervasive modifications to UNIVAC 1100 LISP is the inclusion of the paging environment. Basically, this paging environment allows the user a virtual reference space of 131K while the real core allocated him is only a fraction of this figure. All pointers and references in the LISP system are made with virtual addresses which require translation to real addresses. Both virtual and real space are divided up into equal-sized sections called blocks; these blocks correspond to the LISP notion of page in that only one node type may reside on a block. Virtual blocks can be in one of three states:

(1) Available - Block has not been requested by LISP.

(2) Core Resident - Block has been requested by LISP and has been assigned a core block.

(3) Drum Resident - Block has been core resident at one point, but has relinquished its core block for use by another virtual block. It now resides on a drum file assigned to the LISP system.

Core blocks are either 1) allocated to a virtual block or, 2) available for assignment. To keep track of the various blocks' states, there are four data structures:

(1) Available Virtual Block Queue

(2) Available D-Bank Core Block Queue

(3) Available I-Bank Core Block Queue

(4) Page Table - This structure also records each virtual block's node type and, if the virtual block is core resident, the core block assigned it.

Each virtual block has a unique place for itself in the drum file. If s is the number of sectors that can store one block, then virtual block 0 occupies the first s sectors of the file, virtual block 1 occupies the second s sectors of the file and so on, thus, the drum file may be thought of as a replica of the virtual space.

When LISP requests a virtual block for which no core block is available, another virtual block is chosen to become drum resident to free a core block. This choice is made using the Least Recently Used (LRU) Algorithm discussed in Appendix G. With the LRU Algorithm, the virtual block which was referenced the longest ime ago is the one selected to reside on the durm. The LRU Algorithm was used instead of the Weighting Algorithm (described in Appendix F) primarily for its speed in processing each virtual reference and for its simplicity of implementation.

Certain virtual blocks are "locked" into core (i.e. once requested they must always remain core resident). Most of these virtual blocks are also "fixed" in core. In this case, the virtual block and its assigned core block have the same number so that all virtual addresses in the virtual block are identical to their real addresses. The virtual blocks containing system code and data, system atomic symbols, and register space are all fixed and locked to greatly reduce the number of reference translations in the system. Virtual blocks containging compiled LISP code are defined to be locked to maintain execution efficiency due to the probable high number of references to compiled functions within STIS. However, compiled code blocks need not be locked and may be defined to be pageable if it is seen that there is little activity in these functions.

All of the system functions were amended to accommodate the paging environment. These changes are transparent to the user except for functions such as *EXAM which now use virtual addresses as arguments. Also, a garbage collection is now performed automatically when either 1) no available virtual blocks exist when one is needed or 2) no available I-Bank core blocks exist when one is requested.

A subroutine TRAPPER was written to perform the following functions:

(1) translate a virtual address to a real address

(2) operate the LRU algorithm

(3) write a block onto the drum file

(4) read a block from the drum file, and

(5) keep count of the number of reads and writes to the drum file.

This routine is called within the system code whenever a data item is read from or written to a virtual address. A modified version of this routine was written especially for use by compiled code. The LISP compiler was then modified to output the calls to this routine in the appropriate code generation functions. The compiler was also modified to generate instructions only with virtual address fields.

6.2.2 Permanent Storage Facilities

The usage of LISP in a large data base context requires that a facility exist to provide for the "permanent" storage of completed list structures, so that they can be made available for general future retrieval. Such a facility requires a mass storage I/O and allocation/deallocation scheme that is considerably more elaborate, in the long-term development, then that required for temporary paging.

The LISP system has been modified to provide facilities for the storage, and subsequent access of list structures in a compressed format on external mass storage files.

A list structure to be output onto mass storage must be the value of an atomic symbol identified in the output call. Conversely, a list structure brought in from mass storage becomes the value of an atomic symbol identified in the input call.

To provide these services, three functions have been added to the LISP repertoire:

(1) (PUTNODE atomic symbol block number partition id.)

The value of the specified atomic symbol is output to the specified block location in the specified partition (file).

During output, the list structure is converted to printable form (i.e., to an S-expression) and packed into an output block area. When output characters exceed the capacity of the output area, another block location on external storage is allocated, the allocated block number is placed as a link in the last word of the output area, and the output area is written to the currently specified block location. The newly-allocated block becomes the currently specified block location, and packing of output characters commences at the beginning of the output area. This process continues until the list structure is exhausted, at which time the current contents of the output area are written to the currently specified block location on external storage.

(2) (GETNODE atomic symbol block number partition id.)

Data (in S-expression form) are read beginning at the specified block location on external storage, into an input block area and translated into an internal list structure using existing LISP read functions. If data in the current block are exhausted before the list structure is completed, the block number of the next block on external storage containing continuation data is obtained from the link word in the current block. Blocks are read as required until the entire S-expression has been processed and the list structure has been completed. The resulting list structure is assigned as the value of the specified atomic symbol. A list of any link blocks (continuation blocks) encountered during the processing of the S-expression is assigned to the system atomic symbol SEGLIST. (This list would be required for future "replace" and "release" operations).

(3) (ALLOCATE # Blocks required partition id.)

The specified number of consecutive blocks of external storage is reserved for subsequent use by the caller. The allocate function is used for the acquisition of unused blocks for the initial output of internal lists, and is also used (transparently to the caller) during PUTNODE processing to acquire continuation blocks. The allocation scheme is very simple, but adequate to allow for the building and debugging of higher-level functions. It is anticipated that it will be replaced by a more complex allocation algorithm tailored to system requirements yet to be defined.

6.2.3 Interface for LISP Callers

A facility has been designed to allow the LISP system to be called by non-LISP programs. The calling mechanism provides for two main components:

(1) the passing to LISP of an S-expression which LISP can evaluate in the normal way.

(2) the identification, description and location of value parameters which are transformed from LISP to non-LISP format or vice-versa.

A literal in the calling program designated as the LISP S-expression is transferred by the interface routine to the LISP read buffer. The literal must be a complete expression in the LISP language. The expression is not scanned by the interface routine.

Each value parameter associated with the call is described by means of a 2-word packet. This packet includes the following information:

(1) the name of the parameter - this name is interpreted as the atomic symbol to which the parameter value will be attached (either through the action of the interface routine or through the evaluation of S-expression by LISP).

(2) input/output indicator - indicates whether the parameter value is supplied by the caller, or is to be delivered to the caller after the expression is evaluated.

(3) array indicator - indicates whether the parameter is a single value or a one-dimensional array.

(4) parameter type (e.g., integer, double-precision floating point)

(5) parameter size - for an array, the number of elements; for a string, the size in words (arrays of strings cannot be defined).

After placing the S-expression into the LISP read buffer, the interface routine examines each parameter description packet. The routine assures that an atomic symbol having the name of the parameter is defined. If it is an output parameter, no further processing of it takes place at this time. If it is an input parameter, the parameter value is placed into the type of LISP space appropriate to the parameter type. If an input array is being processed, as each element is assigned LISP space, it is added to a LISP list structure. When the input parameter has been completely processed, its LISP value is assigned to the appropriate atomic symbol.

After all parameters have been processed, the LISP evaluation process is entered.

On normal return from LISP evaluation, the interface routine rescans the parameter description packets for output parameters. For each output parameter, the LISP value for the appropriate atomic symbol is retrieved and placed into the specified user area. During this processing, appropriate checks are made to determine that the LISP value types correspond to those given in the output parameter description. A LISP list must correspond to a defined array. No change is made to LISP values during this processing.

Error-free processing of a user call results in a return to the user with a zero status word (in a location specified in the calling sequence). Errors occurring during parameter processing, or during evaluation, result in a return to the caller with a non-zero status word.

The first call to LISP made by the user results in the initializing of the LISP system.

The above described process has been coded and tested. Complete implementation of the interface, however, still awaits the trapping of all possible evaluation errors, and the development of a contingency processing strategy suitable for the destined environment.

6.2.4 LISP Address Space

Currently, LISP is oriented toward the use of an 18-bit (half word) pointer, of which only 17 bits are available for addressing. This address limit of 131K words is too low to provide for complete addressing over the whole range of the projected STIS data base. In order to avoid a requirement to modify LISP in this regard (though the modification is conceptually simple, it would at best result in a considerable reduction in available core usage), it has been decided that each node in the concept net will be a separate list. Several methods are available for dealing with these lists; the method that seems most attractive currently is to treat the node identification as an atomic symbol and to list the attribute and value structure of the STIS Node

from the atomic symbol. Each list, hence each node, can be up to 131K words long. Nodes which are logically larger than 131K words can be segmented by the user using attributes defined for that purpose.

Several system utility functions were implemented within LISP to help the programmer control the free space utilization of the system.

(1) (REMOB x) - removes all atomic symbols in the list x from the hash list so that they are eligible for garbage collection.

(2) (GC) - allows the programmer to invoke the garbage collector to remove dead space and place it on the appropriate available free-space list.

6.2.5 Data Type Double Precision Real

A new node type was created for double precision reals. Each double precision real is stored as a normalized double precision floating point number. Its I/O format is similar to single precision with a 'D' (in place of an 'E') preceding any exponent. A decimal '.' and a 'D' must be present when inputting the number. The arithmetic routines were modified to convert arguments to double precision and return a double precision number if any arguments are double precision. The following was also implemented:

(1) (FPCOMPRESS DP) - converts the double precision number DP to a single precision value and returns that value.

The I/O routines which perform real number conversions (BCD to binary and vice-versa) were also improved for both speed and accuracy.

6.2.6 Other Modifications

Certain additional modifications were made to LISP due to the impact of one or more of the primary items of investigation mentioned earlier as well as the planned environment within which LISP would function. These modifications include the following:

(1) Compiled code location - Compiled code has been restricted to occupying I-Bank space only.

(2) Dynamic core expansion - The LISP system function (GROW) was removed as it was deemed unnecessary due to the paging implementation.

(3) Control Stack overflow - Tighter security measures were installed to detect control stack overflows into any contigiously collected data areas.

(4) Value stack placement - The placement of the value stack and the LISP data bank (D-Bank) in general was made Collector dependent and assembly independent. A relocatable element of LISP can have its starting D-Bank address specified by a Collector DBank directive. The Collector bank-naming technique is now used to create an absolute element of LISP. Explicit bank collection and control stack overflow traps allow other externally assembled modules to be mapped together with LISP into one executable element.

(5) Timing information - (TIME) and GCTIME) functions were modified to return the total accumulated SUP (Standard Unit of Processing) time for the LISP session and for garbage collections respectively.

The control card :TIME was augmented to provide the total accumulated SUP time for the session, the CPU time, the I/O time, the Executive Request time, and the core block residency time.

(The reverse of this page is blank)

SUPPLEMENTARY BIBLIOGRAPHY ON ASSOCIATIVE PROCESSING

Berra, P.B.: Some problems in associative processor applications to data base management. NCC, 1974. 1-5.

Davis, E.W.: STARAN parallel processor system software. NCC, 1974. 17-22.

DeFiore, C.R. and P.B. Berra: A quantitative analysis of the utilization of associative memories in data management. IEEE Trans Computers. C-23.2 (Feb 1974). 121-133.

Feldman, J.D. and L.C. Fulmer: RADCAP-An operational parallel processing facility. NCC 1974. 7-15.

Kyburg, H.E. and H.E. Smokler: Studies in Subjective Probability. Wyley, 1964.

Lea, R.M.: Information processing with an associative parallel processor. Computer. Nov 1975. 25-32.

Minsky, N.: Rotating storage devices as partially associative memories. FJCC 1972 (AFIPS Vol 41). 587-95.

Moulder, R.: An implementation of a data management system on an associative processor. NCC 1973. 171-76.

Mutt, G.J.: An overview of a multi associative processor study. ACM Proc. 1974, Vol. 1. 101-104.

Ozkarahan, E.A., et.al.: RAP - An associative processor for data base management. NCC 1975. 379-87.

Prentice, B.W., et.al.: Associative Processor Application Study. RADC-TR-74-326 (Jan 1975), (A005308).

Rudolph, J.A. et.al.: The comming of age of the associative processor. Electronics, Feb. 15, 1971.

Summers, M.W.: An Associative Processor Application Study. RADC-TR-75-318 (Jan 1976), (A021232).

Thurber, K.J. and L.D. Wald: Associative and parallel processors. ACM Computing Surveys. 7.4 (Dec 1975). 215-53.

Turn, R.: Computers in the 1980's. Columbia University Press., 1974

BIBLIOGRAPHY

Carnap, Rudolf and Jeffrey, Richard C., Studies in Inductive Logic and Probability, Vol. 1, Univ. of California Press, Berkeley, Ca., 1971.

Duda, R.O., et al, "Subjective Bayesian Methods for Rule-Based Inference Systems." NCC, June 1976.

Fishman, Daniel H. et al, "MRPPS--An Interactive Refutation Proof Procedure System for Question Answering," Univ. of Maryland Computer Science Center, College Park, Md., February 1973.

Fishman, Daniel H. et al, "The Q* Algorithm--A Search Strategy for a Deductive Question-Answering System," University of Maryland Computer Science Center, College Park, Md., February 1973.

Gettys, Charles F., and Willke, T.A., "The Application of Bayes' Theorem When the True Data State is Uncertain," Organizational Behavior and Human Performance, 4, (1969) pp 125-141.

Goldhirsh, I. and R. Carson: "A Deductive System for Intelligence Analysis" AAI-2255-TR-1, April 1, 1976.

Johnson, Edgar M., "Numerical Encoding of Qualitative Expressions of Uncertainty," Army Research Institute for the Behavioral and Social Sciences, Technical Paper 250, December 1973.

Johnson, Edgar M., "The Effect of Data Source Reliability on Intuitive Inference," Army Research Institute for the Behavioral and Social Sciences, Technical Paper 251, July 1974.

Keuhner, D., "Some Special Purpose Resolution Systems," Meltzar and Mechie (eds.), Machine Intelligence, Vol. 7, Wiley (1972), pp 117-128.

Kling, Rob., Fuzzy Planner--"Computing Inexactness in a Procedural Problem-Solving Language," Technical Report No. 168, February 1973.

Kowalski, R., "And-or Graphs, Theorem-proving Graphs and Bi-directional Search." Meltzar and Mechie (eds.), Machine Intelligence, Vol. 7, Wiley (1972), pp 167-194.

Kuhns, J.L., "Probability Appraisals Based on Possibly Unreliable Reports," Operating Systems, Inc., OSI Technical Note No. 9, OSI:N73-004, April 5, 1973.

Lee, Richard C.T., "Fuzzy Logic and the Resolution Principle," Journal of the Association for Computing Machinery, Vol. 19, No. 1, January 1972, pp 109-119.

LeFaivre, Rick, "Fuzzy: A Programming Language for Fuzzy Problem-solving," National Technical Information Service, PB-231 813, January 1974.

BIBLIOGRAPHY (Continued)

Levine, Jerrold M. and Eldredge, Donald, "Effects of Ancillary Information Upon Photointerpreter Performance," Army Research Institute for the Behavioral and Social Sciences, Technical Paper 255, September 1974.

Sable, J., "Design Concept for an Augmented Relational Intelligence Analysis System (ARIAS)," AUER-2022-TR-2, August 1973, RADC-TR-73-342, (773189).

Shapiro, Bruce A., "A Survey of Problem Solving Languages and Systems," National Science Foundation, Technical Report TR-235, March, 1973.

VanderBrug, Gordon J., and Minker, Jack, "State-Space, Problem-Reduction, and Theorem Proving--Some Relationships," Communications of the ACM, Vol. 18, No. 2, February 1975, pp 107-115.

Zadeh, L.A., "Outline of a New Approach to the Analysis of Complex Systems and Decision Processes," IEEE Transactions, Systems, Man, and Cybernetics, Vol. SMC-3, No. 1, January 1973, pp 28-44.

(The reverse of this page is blank)

APPENDIX A

THE CONCEPT NET -- A NEW INFORMATION STRUCTURE FOR STIS

1. A MODEL FOR INTELLIGENCE INFORMATION

The Concept Net represents the intelligence analysts' collective view of the current state-of-affairs in the real world. It is populated with "facts" distilled by the analyst from observations, reports of observations, and assertions concerning his sphere of interest in that world. Since it is a dynamic world and viewed, as it were, "through a glass, darkly", each fact has associated with it an open-ended set of qualifying statements which include, typically, the source (or message) from which it was derived, its interval of validity (in time and/or space), the date of observation or entry into the system, the credibility (probabilistic truth value) assigned by the analyst, and the time-constant (or "half-life") which characterizes the volatility of the information. Because virtually all intelligence information is both of questionable veracity and subject to change, we view the original credibility level as being modulated by an exponentially decaying weighting function whose time constant is characteristic of the volatility of the type of information in question. For example, the place of employment of an individual may have a half-life of four years. That is, if it was reported in 1970 that George Murphy worked for RCA, and this "fact" was accepted with a credibility of 0.8, then in 1974, in the absence of any new data concerning Mr. Murphy, the credibility of that fact would be 0.4.

Another characteristic of intelligence information is that there may be conflicting reports concerning the facts about a given entity and/or legitimately differing views among one or more analysts as to what the facts may be, or, for that matter, more than one value for a given attribute may be valid in a given time interval. (The case may be that Mr. Murphy, while working for RCA, moonlights as an instructor for Rutgers University so that apparently conflicting reports on Mr. Murphy's occupation may be reconcilable, and coexist with a high credibility. On the other hand, the report of Mr. Murphy's employment at Rutgers may be a deliberate plant or "cover" to obscure the fact that he works for RCA.) For this reason, and to provide for simultaneous use of a common body of information among many analysts, the analyst or organization which is responsible for a given "fact" is recorded as part of the information qualifying that fact.

2. THE CONCEPT NODE

The Concept Net is organized as a network of nodes, each of which represents a concept (such as an individual or other entity) which is of interest to the analyst. The node in turn contains a set of facts (properties) made up of attribute names and qualified values, which describe the entity and its relationships to other entities. These facts are derived from (and tied to) messages concerning observations of the real-world. Other nodes may represent concepts which exist independently of messages (or observations of the world) such as semantic concepts representing the attributes and values themselves, as well as their inter-relationships. (Value nodes will also be related to the entities which are described by (or use) those values, providing a cross-index to the Entity Net.)

Each node in the Concept Net comprises an open-ended set of properties of the concept or real-world entity which is represented by the node. A property is an attribute-name/attribute-value pair which may, in turn, be qualified by an arbitrary list of properties. Attributes and values (also terms and words) are themselves represented by nodes in the Concept Net. An entity node may stand for a real-world individual, unit, facility, weapon, event, etc. A node may also represent a state or sub-entity attached to a parent node. For example, a parent node may represent a generic class of weapons, such as the Minuteman missile, while a sub-node may represent a specific example of that missile installed at a particular site, with a particular target, etc.

When a given entity or other concept node (the source) bears some relationship to another concept node (the target), that relationship is represented in what is called an entity-relational attribute in the source node. Its value is the identifier (Node #) of the target node. In order to provide complete cross-referencing, there will be defined for each relational attribute R (using its Attribute Node in the Semantic Net) an inverse relation R^{-1} so that if entity $\underline{a}$ bears relation R to entity $\underline{b}$ "R(a b)" then entity $\underline{b}$ bears relationship $\underline{R}^{-1}$ to entity $\underline{a}$ "R^{-1}(b a)". For example, if the Pershing Missile has a test site at White Sands "Has Test-site (Pershing Missile, White Sands)", then White Sands is the test site of the Pershing Missile "Is test-site (White Sands, Pershing Missile)" where "Inverse (has test site, is test site)"

and "Inverse (is test site, has test site)". In the above example, the first two statements would be in the Entity Net (Pershing Missile and White Sands nodes, respectively) while the latter two statements would be in the Semantic Net (Has test site and Is test site nodes, respectively).

In addition to entity-relational attributes, an entity may possess attributes whose values are names, numbers, or descriptive terms which are not other entities. These values may be represented by nodes in the Semantic Net (rather than the Entity Net) which in turn cross reference, as entity (or index) lists, those entities which use them. Hence, the distinction between entity-relational and non-entity-relational attributes has little operational significance for search strategies in the system. In either case, the entities possessing a given property are accessible through the cross-referencing (indexing) feature, whether it be the node representing the target of an entity-relational attribute or the node (in the Semantic Net) representing the value of a non-entity-relational attribute. The entity list under the value node can be considered the inverse of the non-entity relational attribute in the entity node in which it occurs. The Concept Net provides for both an attributes-under-entity (normal file) and an entities-under-attribute (inverted file) point of view. This redundancy of access path -- sacrificing space for time -- is built by the system, under control of the Data Base Administrator (who may limit this redundancy selectively) and need not concern the analyst who chooses to limit his role to that of an information consumer.

3. SUB-NODES -- COMPOSITE ATTRIBUTES AND N-TUPLES

There will be instances in the Concept Net when it will be useful to consider one node as subordinate to another in a hierarchic sense (rather than the non-hierarchic, or coordinate, relationship between two nodes which are joined by an entity-relational attribute). When this subordinate relationship is defined, it implies the desirability to store the subordinate node so that it is physically accessible with the parent node, reflecting logical dependency and/or predictable access patterns. When this occurs, the subordinate node is called a sub-node of the parent, or master, node.

The sub-node relationship can arise in several contexts. In addition to the close master/slave relationship that may exist between two entities, mentioned above, a subnode may represent what is called a composite attribute, or n-tuple. A composite attribute is an attribute comprising a set (n-tuple) of simpler attributes. For example, position may be defined as a composite attribute comprising the simple attributes latitude and longitude, or address comprising number, street, city, and state. Composite attributes provide for generic terms which conveniently reference and retrieve a set of specific information. The analyst or programmer who is concerned about the structure of the Concept Net or is developing appropriate terminology for semantic concepts may work with the Data Base Administrator to define composite attributes or other sub-node relationships.

4. FORMAL DESCRIPTION OF THE CONCEPT NODE

Each Concept Node in the STIS Concept Net is represented by the same formal structure, called a description list. Entity Nodes, Attribute Nodes, and Value Nodes are all instances of Concept Nodes in STIS. Each Concept Code (or Node #) is the name of a description list. (The Node # will be used as a key to obtain the description list from permanent storage.) Subnodes are also represented by description lists but they do not have separate Concept Codes associated with them since they are stored with the parent node. A composite value (the value of a composite attribute) is a special case of a sub-node in which the attributes have been predefined.

A Concept Node in STIS is a description list. The formal syntax of a description list is specified in Table A-1. Lower case letters represent syntactic variables and upper case letters represent concept codes or other terminal atomic symbols.

There is no syntactic distinction between brackets and parentheses. Note that a description list is defined recursively so that there is no constraint on the nesting of subnodes representing qualifiers or composite values.

TABLE A-1

NODE STRUCTURE SPECIFICATION

Note: The convention used here for syntax specification uses the following metalinguistic symbols:

←	is defined as, or can be replaced by
⸤ ⸥	one or more occurrences of the expression enclosed by the lower half-bracket
|	choice symbol
⸢ ⸣	optional (at most one occurrence) of the expression enclosed by the upper half-brackets

Syntax

deslist ← [@ ⸤prop⸥]

prop ← (A val)

val ← V | [⸤val⸥] | (val qual) | deslist

qual ← [* ⸤prop⸥]

Semantics

deslist	=	description list
prop	=	property
val	=	value
A	=	Attribute Code (i.e., Node #)
V	=	Value Code (i.e., Node #), numeral, or string representing a terminal value.
qual	=	qualification list
[V ...]	=	array (list) value
(val qual)	=	a qualified value; qual is the subnode which qualifies val
(A deslist)	=	a composite property; A is the composite attribute and deslist is the sub-node representing the composite value.

TABLE A-1 (Continued)

Examples

The following description list examples represent Entity Nodes. In the interest of clarity, attribute and value names are used rather than Node Numbers.

#1 = [@(Name [@(First Jerry) (Last Sable)]) (Age (45 [*(Source Est)
(Accuracy ± 3) (Validity-interval 1975)]))
(Works-at #2)]

#2 = [@(Name AAI) (Fac-type Consultant-org) (Employs [#1 #3 #4])
(Location ([Phila Wash] [*(Cred 0.90]))]

#3 = ([@(Name Schernecke) (Works-at #2)] [*(AOR Consultants)])

#4 = ([@(Name McCrea) (Works-at(#2[*(Validity-interval [1963 1975])]))]
[*(AOR Consultants)])

5. THE SEMANTIC NET

The Semantic Net is that subset of the Concept Net comprising Attribute and Value Nodes. All attributes and values are represented by nodes in the Semantic Net. (In the case of numeric values, the Value Node represents an interval on a lograrithmic or linear scale.) When they occur in the description list of a node in the Entity Net, attributes and non-numeric values are represented by their Concept Codes (Node Numbers).

5.1 Attribute Nodes

The description of any concept consists of a list of properties, i.e., attribute name/value pairs. Since attributes are concepts themselves, they are represented by nodes in the Semantic Net subset of the Concept Net. Some of the attributes which can be expected to be used in the description list of an Attribute Node are listed below. (It should be noted that as in all nodes, these attributes, except when they are self-referencing, are represented by the Node Codes of Attribute Nodes. Their values are represented either by Node Codes or by Term Codes.)

Attribute name

Synonyms

Narrower attributes (for composite attributes)

Broader attributes (for components of composite attributes)

Inverse attribute (for Entity Relational attributes)

Values (the list of values for this attribute, limited to the first domain element in the case of Entity Relational attributes)

Attribute Data Information -- the value of this attribute is a pointer to the Attribute Data Record in a Direct Access file outside of the Concept Net. The ADR defines the format, precision, units, and "owner" of the attribute. This is an example of a special attribute, or Process Hook, which invokes an outside routine to compute a complex value, using the nominal attribute value as a parameter.

Other Attribute properties, such as transitivity, reflexivity, and symmetry which may exist will also be represented in the property list of the Attribute Node.

5.2 Value Nodes

Each non-numeric value, or range of numeric values, which can serve as a retrieval condition will be represented as a Value Node in the Concept Net. When indicated by the analyst, or Data Administrator, the Value Node will serve as the head of an index to information in the Entity Net. This provides support for the three basic strategies for retrieving information about intelligence entities:

(1) through the context of an explicitly identified entity, including its association with other entities via relational attributes,

(2) through a retrieval criterion made up of a set of specified properties which the entity should possess, and

(3) through properties which are plausible for the entity because they can be inferred from generalized rules stored in the Concept Net.

Some of the attributes which can be expected to be used in the description list of a Value Node are listed below:

Value name

Synonyms

Narrower values (or subsets)

Broader values (or supersets)

Attribute (the attribute that has this node as a value, the inverse of the Values attribute in the Attribute Node)

Entities (the entities which have this node as a value. This serves as the index list for those entities.)

6. THE ENTITY NET

Information about any intelligence entity of concern to the analyst can be stored in STIS by creating an Entity Node to represent it in the Concept Net. Once the node is created, the description of the entity is stored as a list of properties. Internally, the entity is known by its Node Number, which serves as its retrieval key from permanent storage, as is the case for any node in the Concept Net. In its simplest form, the entity number n is represented by a description list such as:

n = [@ (A a) (B b) (C c) ...]

The interpretation is that the entity represented by node n has all of the properties listed. That is, in conventional relational or logical format, the attributes A, B, C, ... are binary relations connecting the entity and a value and the following conjunction holds:

$$A(n,a) \wedge B(n,b) \wedge C(n,c) \wedge \ldots$$

Thus, in the Entity Net, information is collected in an "attributes-under-entity" format, while in the Semantic Net, one may say that the same information appears in an "entities-under-attribute" format. As will be discussed below, the simple description list form can be generalized in a number of important ways.

6.1 Entity Relations

The simplest relations are attributes which take scalar values, either literal or numeric, such as Name(n,Atlas) and Weight(n,150). However, values are generalized to permit arrays, such as Name(n,[Atlas,M12]) and Location(n,[ND,FL]). Assuming Node n is #10, this would appear in description list format as:

#10 = [@(Name [Atlas M12]) (Weight 150) (Location [ND FL])]

Entity-relational attributes name other Entity Nodes as values. If entities #11 and #12 were test sites for #10, then the update command "Add Test-Site (#10, [#11, #12])" would add the property (Test-Site [#11 #12]) to the description list for #10.

By permitting a value to be represented by a description list, or subnode, the descriptive power of the system is augmented in a number of ways. The simplest instance of this, the composite attribute, was described in Section 3. Other cases will be discussed in the following paragraphs.

6.2 Generic Entities

It is often useful to describe an object as a generic type for which, in the real world, there exists a number of specific occurrances. This can be done by creating a node, called a generic entity, which represents the common characteristics for these objects. This can then be supplemented by a node for each individual object for which specific information is required but which is not characteristic of the class as a whole. For example, suppose we have the missile type Atlas represented by:

#20 = [@(System ICBM) (Name Atlas) (Weight 150)
(Accuracy 3) (Instances [#21 #22 #23])]

Nodes #21, #22, and #23 then are specific entities whose general characteristics are given in node #20 and therefore may be inferred by reference and need not be explicitly repeated. Each instance will reference the generic entity and give only unique characteristics, such as:

#21 = [@(Location ND) (Target #31) (Serial 1234) (Generic-entity #20)]

Note that the attributes "Instances" and "Generic-entity" are a converse pair.

6.3 Entity States

It is often necessary to track changes in a given set of properties of a specific object. To do this, subnodes called "states" are created.

The relationship between a specific entity and a state of that entity is parallel to that between a generic entity and a specific entity. That is, only properties whose values change from one state to the next need be recorded. Invarient properties are given in the parent node. For example, suppose a Polaris type submarine is being tracked. Intermittent reports of its location may be given in state nodes which reference the specific entity node. The specific entity node may, in turn, reference a generic entity. This interrelationship of subnodes is diagrammed in Figure A-1.

The recurring motive for introducing subnode relationships such as "instance" and "state" is to avoid redundant storage of information. The payoff for eliminating unnecessary redundancy is reduction of maintenance and retrieval time as well as space. Storage compression at the state level can be carried to a further stage when changes in state are predictable or can be represented analytically as a function of time. Opportunities for this may exist in situations such as when a periodic itinerary for a submarine or other ship is known, or when a satellite position may be found from orbital parameters rather than extrapolation or interpolation of tracking data. In such cases, state nodes may be replaced by compact state-transition information.

6.4 Fact Qualification

It is possible to modify or qualify information by appending a qualification list to either a description list (node or subnode) or a value. The qualification list has the format of a description list so that the two forms are respectively (deslist qual) and (val qual) where the second element is the qualification list. Typically, qualification information in an Entity Node will contain fact control (access control) information if it is at the node level and fact control and/or source, credibility, and temporal data at the value level. Because information may be obtained from several sources and may be varying with time, multiple values will be common in the Entity Net. The particular values which are valid for a given analyst at a given time will be determined on the basis of the qualification list.

The default interpretation of the property (A v) for an entity (say e) is that the entity has the value v for the attribute A. In symbols A(e) = v. The value v may either be a scalar V or an array [V...]. However, there

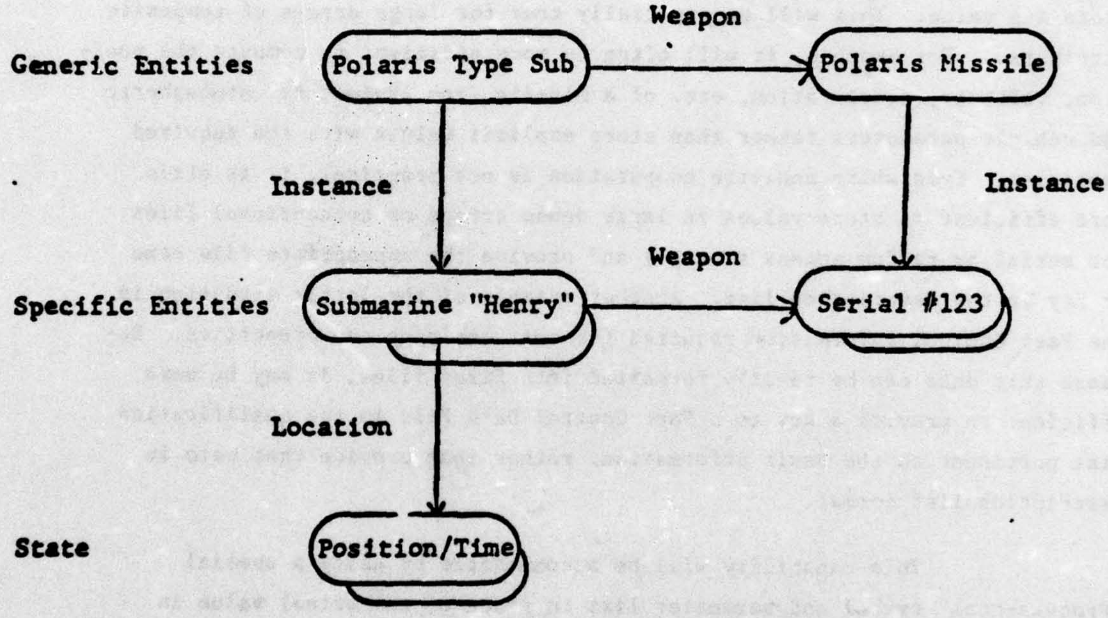

Figure A-1
Subnode Relationships

are occasions when one wants to specify a relational operator other than equality between the attribute and the value. Possible relations are greater-than, less-then, not-equal, approximately-equal, not-greater-than, etc. The qualification list is also the mechanism for accomplishing this, with the exception operator attribute "Rel-op". For example, Age(e)>40 would be given as (Age(40[*(Rel-op >)])).

6.5 Computed Values

There will be instances when it is more convenient to compute a value for a given attribute from specified parameters rather than explicitly store its value. This will be especially true for large arrays of composite attributes. For example, it will often be more efficient to compute the position, velocity, acceleration, etc. of a missile from trajectory, atmospheric and vehicle parameters rather than store explicit values with the required precision. Even where analytic computation is not practical, it is often more efficient to store values in large dense arrays or conventional files (on serial or random access storage) and provide the appropriate file name or key in the description list. Another example of the latter situation is the Fact Control Information required for most entities and properties. Because this data can be readily formatted into fixed files, it may be more efficient to provide a key to a Fact Control Data File in the qualification list pertinent to the basic information, rather than provide that data in description list format.

This capability will be accommodated by using a special "Process-Hook" symbol and parameter list in place of the actual value in the description list. The retrieval mechanism, when encountering the Process Hook, will invoke the specified program and supply the given parameters. The called program will return the required value.

6.6 Quasi-transitive Relationships

The use of entity relational attributes in the description of the various objects of interest to the analyst results in a network of nodes in which information is highly associated. This richness of association

permits information to be retrieved from many points of view or search paths. Although this feature is, in general, desirable, unless special precautions are observed, there are situations in which it can lead to the retrieval of information which does not validly meet the conditions specified by the interrogator.

Consider, for example, a situation in which a weapon platform (say a fighter-bomber) can be equipped to bear either of two types of armament (say torpedo or incendiaries) depending upon under which service unit (aircraft carrier or tactical air base) it is employed. A given entity relational attribute (such as "uses") may be used to enter this information:

$$\begin{cases} \text{Uses (Carrier Lexington, F-11)} \\ \text{Uses (F-11, torpedos)} \end{cases}$$

$$\begin{cases} \text{Uses (TAC Base Charlie, F-11)} \\ \text{Uses (F-11, incendiaries)} \end{cases}$$

The five entities would then be interconnected with the "Uses" relation as shown in Figure A-2. It is apparent that a request for armament used by the Carrier Lexington (or TAC Base Charlie) may come up with the erroneous answer "incendiaries and torpedos". The fallacy is caused by what can be called a "connection trap" in the F-11 "hub" of the network. It is avoided by using one or both of the following devices:

(1) The set of values of a multivalued attribute are qualified to inform the system that only one of the values can occur in each instance.

(2) A configuration node (or subnode) is created to describe each valid configuration of properties.

These approaches are detailed below.

Since an attribute may have an array as a value, we can have a property such as:

(Armament [torpedo incendiary])

in the description list for an entity (say F-11). This raises the question as

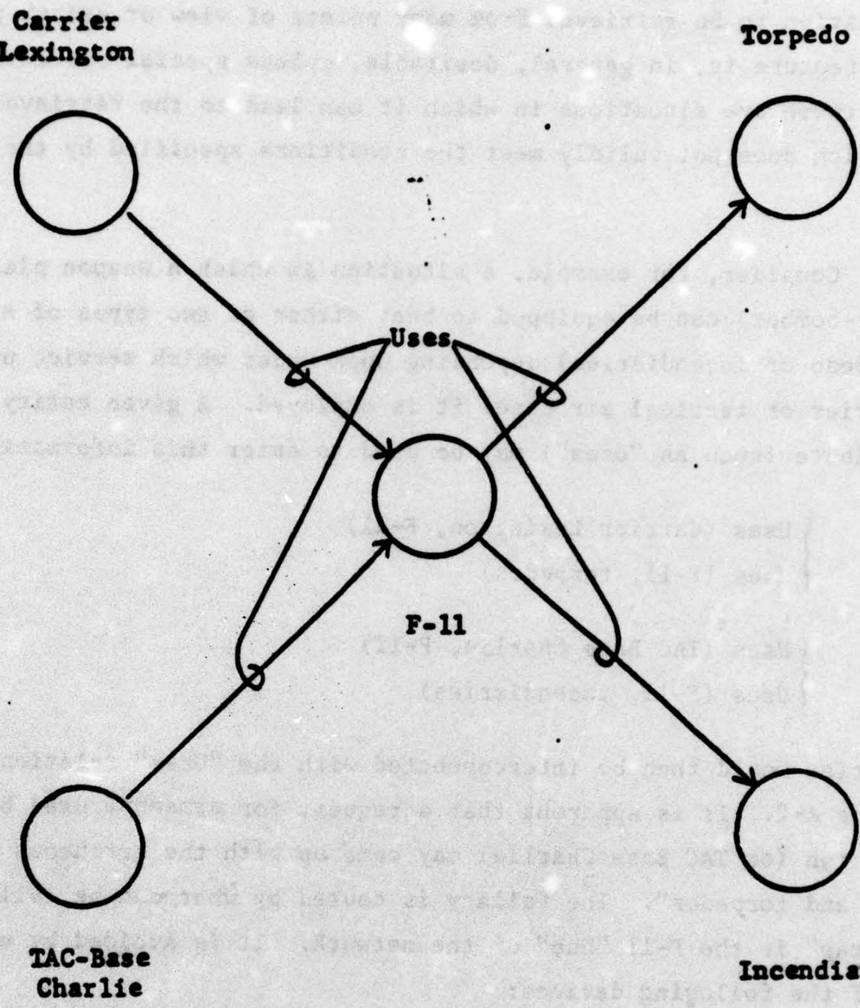

Figure A-2

Connection Trap

to the interpretation of the array:

$$v = [\ v_1\ v_2\ \ldots\ v_n\ ]$$

when it occurs as a value. The members v_i may be an ordered n-tuple, an (unordered) set, a bag (unordered set in which repetitions are permitted), a disjunctive set (any subset is valid), a conjunctive set (all values co-occur), or a choice set (only one value is valid in each instance). The type of set which is intended can be identified by using the attribute "Set-type" in a qualification list for the value. For example:

(Armament ([torpedo incendiary] [* (set-type choice)]))

The use of the "Set-type choice" qualifier alerts the system (and the user) that only one value is valid but in itself is not sufficient to specify which is the valid value in a specific case. This problem can be solved by using a subnode (or a state) of the entity to establish a description of each configuration of the parent entity. For example, we can have the states

[@(Used-by Carrier-Lexington) (Armament torpedo)]

and

[@(Used-by TAC Base Charlie) (Armament incendiary)]

under the generic entity for the F-11. Note that this second approach avoids the multiple values attribute and is sufficient in itself to unambiguously describe the situation.

6.7 Footnotes

The analyst entering facts into the Entity Net will be permitted to qualify any value (or entity) with unformatted comments, warnings, or other text. He simply labels this text (generically called footnotes) with the appropriate attribute (Comment, Warning, etc.) and enters it with other qualification information. Rather than store unstructured text as part of the node, a special use will be made of the Process Hook capability. The value of the specified attribute will be a pointer to the appropriate record in an external Foot note File. The footnote will be retrieved automatically with other qualification information whenever required.

(The reverse of this page is blank)

APPENDIX B

THEORETICAL FOUNDATIONS: THE PROBABILITIES OF COMPOUND PROPOSITIONS

THEORETICAL FOUNDATIONS: THE PROBABILITIES OF COMPOUND PROPOSITIONS

1. INTRODUCTION

A basic consideration in any work with probabilities (or credibilities) of statements of information (or rules) is how to assign a probability to a statement derived from other statements. If the truth-values of X and Y are always 0 and 1, then the values for such compound propositions as $X \wedge \neg Y$ or $X \Leftrightarrow Y$ are clearly defined and well known. But when the range of truth values opens up to the whole interval from 0 to 1, there arises a whole spectrum of possibilities for each statement.

This paper discusses all those possibilities in order to put into proper perspective what has sometimes been given very brief treatment in literature on the subject. Some comments will be made on what appears to be a very popular first choice of approaches to the subject. It will also be seen that the second choice, which is usually mentioned briefly, is not the only alternative, although it may be the best.

2. THE DOMAIN

A function of n two-valued variables has a domain of 2^n discrete points, which can be conveniently thought of as the vertices of an n-dimensional orthotope (square, cube, tesseract, etc.). The range of the function may be superimposed as one more dimension. The function or dependent variable may also be two-valued, but is not necessarily so. An example of a mixture of two-valued inputs with multi-valued outputs is the bell-shaped "curve" showing the distribution of numbers of "heads" when n coins are tossed.

The admission of values between 0 and 1 for independent variables literally opens up a whole new world of possibilities. The domain of a function then includes the whole interior and boundary of the orthotope. The knowledge of what goes on at the vertices has been thoroughly investigated and copiously documented in the last few decades. In many cases, these isolated values suggest

what may happen in the rest of the domain, but they never define it with certainty. This will be amply illustrated in this paper, where very different functions defined over the interior of a square will be seen to have identical effects at the corners, and often around the whole perimeter.

The square to be used represents all pairs of values from 0 to 1, inclusive, for two variables, x and y. In the illustrations, this square is shown as if lying flat, as in Figure 0. Vertical coordinates represent values

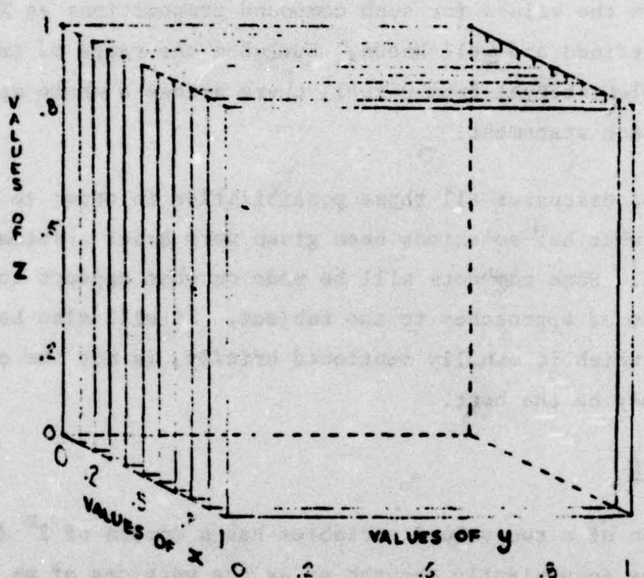

Figure 0. Universe of Discourse

of the dependent variable z. Planes for x = .1, .2, .3, etc. are shown to assist three-dimensional visualization. The planes are drawn as if opaque, but the cube or other solid containing them as transparent. The graphs are in effect continuous truth tables for logical operations on propositions. Capital letters will be used to represent such statements, but small letters represent <u>numbers</u>. If Z = f(X, Y), then x represents the probability that X is true, y the probability that Y is true and z the probability that the compound proposition Z is true.

3. $X \wedge Y$

3.1 Independent

The first function to be considered is $Z = X \wedge Y$. If X and Y are independent, the value assigned to the probability that they are both true is the product of the individual probabilities: $z = xy$.

This is represented by the warped surface illustrated in Figure 1. The lines drawn in the surface represent its intersection with the planes

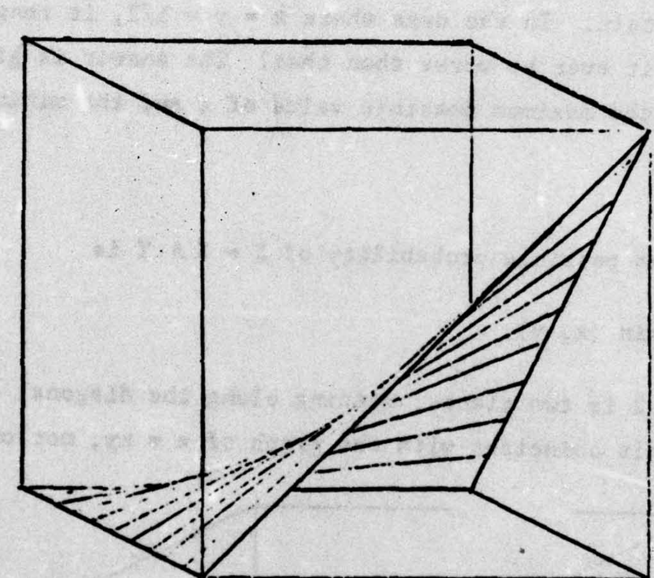

Figure 1. Probability of $X \wedge Y$

$x = .1$, $x = .2$, etc. In this case the intersections happen to be straight lines. The intersections of this same surface with horizontal planes, representing constant values of z, would be hyperbolas. They would display the shape of the surface the way contour lines do on a topographical map.

The independence of X and Y is a very important consideration. If the median height of a population is 1.7 meters, then the statement that a randomly selected member is over 1.7 meters tall has a probability 1/2 of being true. It may be that the same population is equally divided between the sexes.

But the statement

(m is tall) $\wedge$ (m is male)

is likely to be true for more than one quarter of the population. And the probability of a pair of statements, each 1/2 true, can drop quite low, e.g.

(m is male) $\wedge$ (m has ovaries).

In such cases, it would always be helpful (although possibly unconstitutional) to have information that discriminates on all sides of set boundaries. But if the only inputs available are the probabilities of X and Y, then the probability of Z is rather uncertain. In the case where $x = y = 1/2$, it ranges all the way from 0 to 1/2. Can it ever be worse than that? The answer is given in Figures 2 and 3, which show the maximum possible value of z and the minimum, respectively.

3.2 Maximum

The maximum possible probability of $Z = X \wedge Y$ is

$$z = \min(x, y).$$

Its graph in Figure 2 is two planes, meeting along the diagonal of the cube where $x = y = z$. This coincides with the graph of $z = xy$, not only

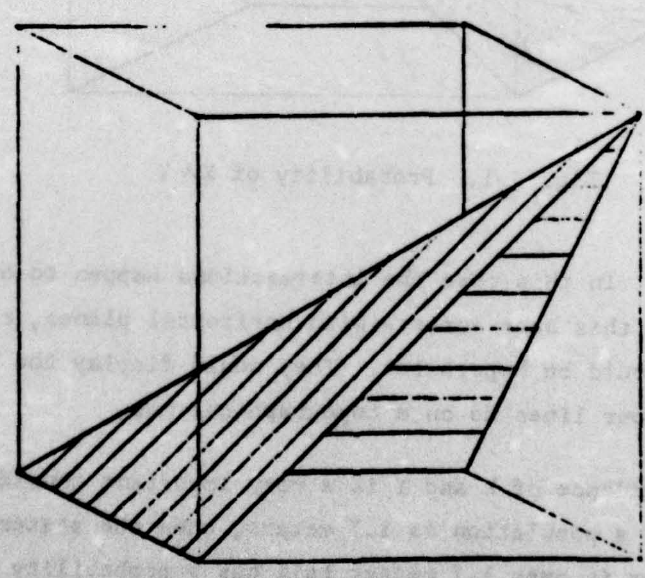

Figure 2. Maximum for $X \wedge Y$

at the four corners of the square, representing the classical truth table, but also along all four sides. This shows that the minimum and the product give identical results if *either* x *or* y takes on a value of 0 or 1. If, for example, X is a tautology, or universally true statement, then the probability of $X \wedge Y$ is simply the probability of Y, as shown by the diagonal line from $y = z = 0$ to $y = z = 1$ in the front face of the cube.

This shows vividly that for the propositional calculus *both* multiplication *and* minimum give the same results. Either xy or min (x, y) gives the truth-table values in the corners of the square. The distinction between them is meaningful only in the interior.

3.3 Minimum

The third function to be graphed in connection with $X \wedge Y$ is the minimum possible value of z. The probability of $X \wedge Y$ increases according to how much the statements X and Y tend to apply to the same set of subjects. The maximum is achieved when one set is a subset of the other. Similarly, low probabilities arise to the extent that the characteristic sets of X and Y avoid each other. If x and y are on the low side, it is possible for the sets to be disjoint, and then the probability is zero. But for larger probabilities the sets become so large that they cannot help overlapping. If they still stay as far away from each other as possible, the territories in which they are *false* are disjoint. The probability then achieved is $x + y - 1$, and it can be done if and only if $x + y > 1$. So the minimum probability of $X \wedge Y$ is

$$\max (0, x + y - 1).$$

This is graphed in Figure 3, where it is seen that it too has the same border as the other two graphs.

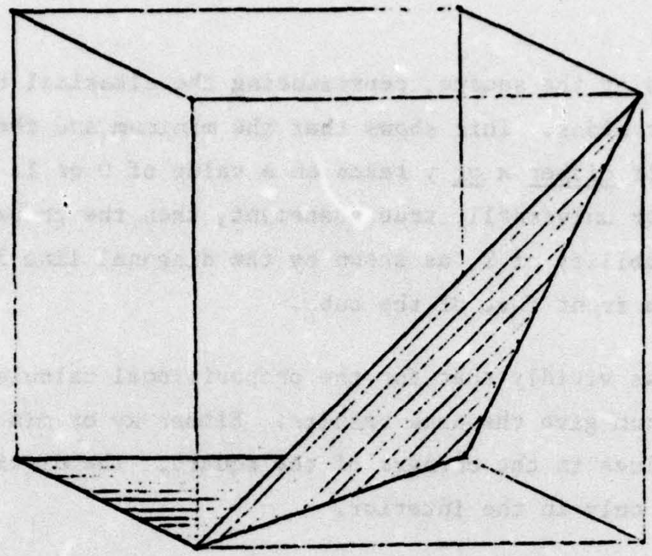

Figure 3. Minimum for X∧Y

3.4 COMPARISON

The three graphs are shown together in Figure 4. In order to show all three surfaces, the picture shows the planes for X = .1, .2, etc. Each one

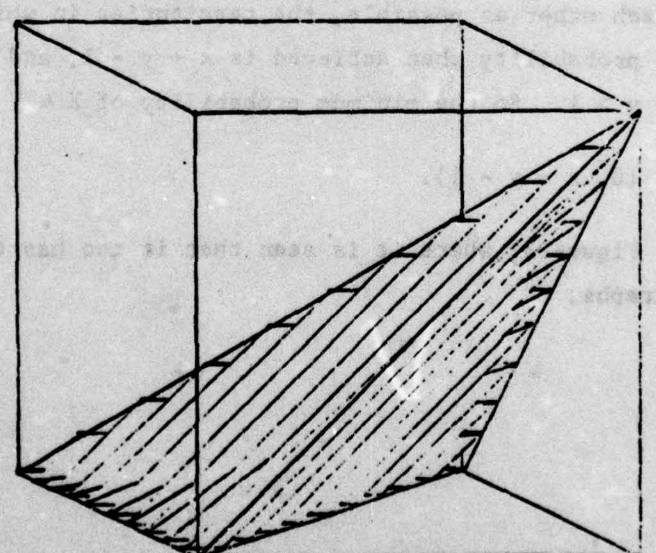

Figure 4. Range for X∧Y

shows the range of z as a parallelogram, which degenerates to a line segment when X = 0 or 1. The product xy is necessarily between the maximum and the minimum, and is seen in Figure 4 as a diagonal of each parallelogram. It is not generally midway between the top and bottom, but it represents an average of all values weighted by their probabilities (assuming X and Y are independent). For example, when x and y are both small, the characteristic sets of X and Y are more likely than not to be disjoint. So the graph of z = xy stays close to the floor in that region, not up near the roof.

The three graphs coincide along all four edges, and it appears as if they differ from each other increasingly toward the center where x = y = 1/2. This is confirmed in Figure 5, where the maximum and minimum are shown as

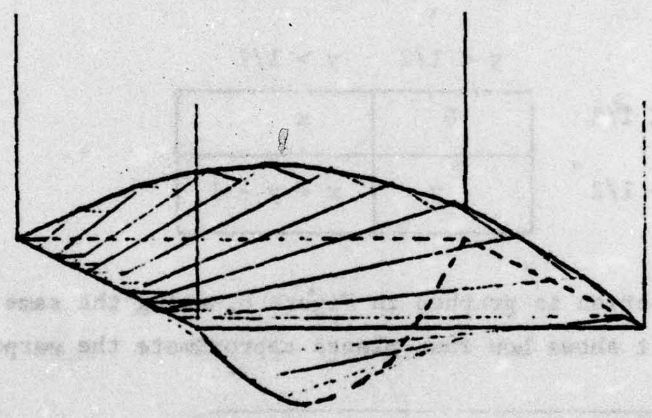

Figure 5. Deviations for $X \wedge Y$

deviations above or below the product xy. This is accomplished by pushing the curved surface down to the floor (z = 0) without changing any vertical distances. Each parallelogram "racks" to a new shape; it is still a parallelogram, but its diagonal is horizontal.

A satisfying feature of Figure 5 is that it shows the minimum for s just as far below xy as the maximum is above xy, but distributed differently. The parabolic ridges and the "tents" they subtend are congruent, but rotated 90^{o} from each other. The planes y = .1, .2, etc. would also intersect these surfaces in parallelograms. Both Figure 4 and Figure 5 show that the maximum probability s = min (x, y) is closest to the probabilistic product at the two corners where x = 0 and y = 1 or vice versa. In the region where both X and Y have high probability (or both low) the minimum function makes an intuitively better approximation of reality.

If someone is determined to use the functions represented by planes rather than the warped surface, it might be reasonable to select the floor or roof according to which is closer. This divides the function into four regions separated by the lines x = 1/2 and y = 1/2.

	$y < 1/2$	$y > 1/2$
$x < 1/2$	0	x
$x > 1/2$	y	x + y - 1

This composite function is graphed in Figure 6, using the same style as Figures 1 to 3. It shows how four planes approximate the warped shape nicely

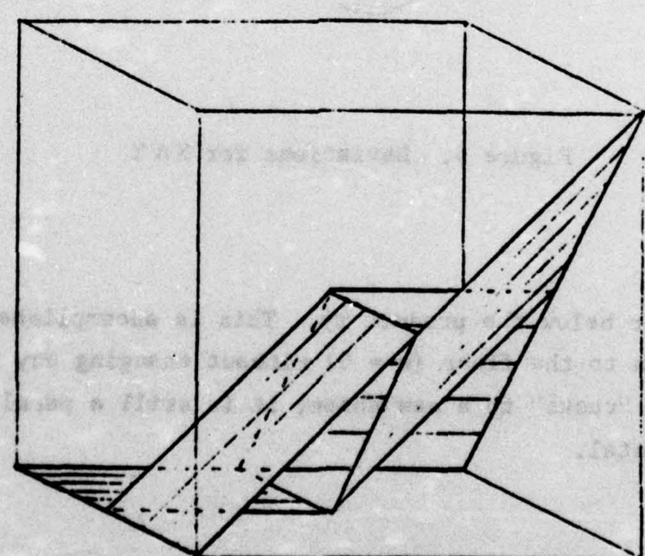

Figure 6. Min or Max for $X \wedge Y$

at the corners, but not very well at the center. This function would in most cases be far less easy to implement then the product xy, but it serves to illustrate an important fact. Even this outlandish function has the same outline as the other three. Thus it would be consistent with the propositional calculus to use the table above to find the probability of $X \wedge Y$. The graph also shows the indeterminacy of this function along the borders between the four quadrants. When $x = 1/2$ the whole parallelogram is still available as a range for z, including the widest possible variation (from $z = 0$ to $z = 1/2$) when $y = 1/2$ also.

4. $X \vee Y$

A long story can be made quite short regarding $Z = X \vee Y$. The maximum probability is $z = \min(x + y, 1)$ achieved when X and Y avoid each other as far as possible. The minimum is $z = \max(x, y)$, achieved when one characteristic set includes the other. The probabilistic formula that strikes a mean between these two, assuming that the statements are independent, is $z = x + y - xy$. The arguments for these assertions are interesting and useful, but Figure 7 shows that the results are simply a variation of those for $Z = X \wedge Y$. The

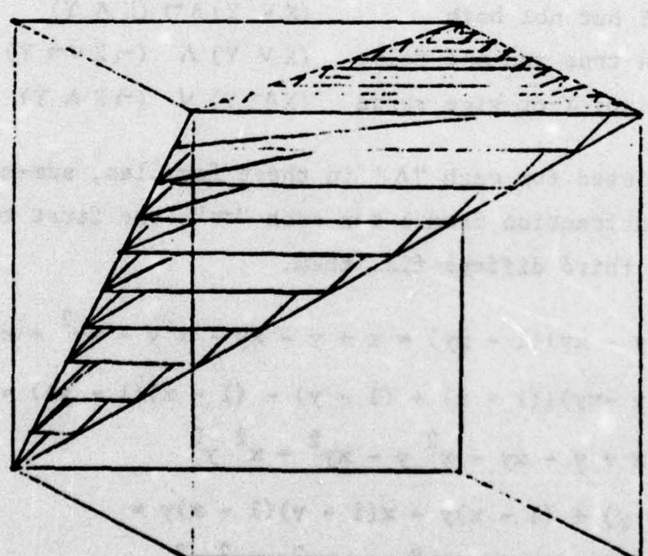

Figure 7. $X \vee Y$

tetrahedron with the warped surface inside is the earlier one turned upside down, reflecting the fact that $X \vee Y = \neg(\neg X \wedge \neg Y)$. This is a variation on deMorgan's Law, and is typical of the conversions that logicians use to make one problem look like another. Such conversions are not without pitfalls, as will be shown in what follows.

5. $X \leftrightarrow Y$

5.1 Independent

The function "exclusive or" is very different from the "inclusive or". Logicians dispatch the difference by pointing out the single "corner" $X \wedge Y$, where one of them takes the value 0, the other 1. The difference is exaggerated when consideration is given to the whole continuum of values for x and y. In applications involving information from various sources, the word "or" may be used for either one, and so blur the distinction. A comparison of the two functions shows how big a mistake it is to confuse one with the other.

Three possible interpretations of "exclusive or" are suggested by English paraphrases that can be used to explain what it means:

X or Y but not both	$(X \vee Y) \wedge \neg (X \wedge Y)$
One is true and one false	$(X \vee Y) \wedge (\neg X \vee \neg Y)$
X and not Y or vice versa	$(X \wedge \neg Y) \vee (\neg X \wedge Y)$

If multiplication is used for each "$\wedge$" in these formulas, sum-minus-product for each "$\vee$", and subtraction from 1 for each "$\neg$", the first two give the same result, but the third differs from them.

$$(x + y - xy)(1 - xy) = x + y - xy - x^2y - xy^2 + x^2y^2$$

$$(x + y - xy)((1 - x) + (1 - y) - (1 - x)(1 - y)) =$$
$$x + y - xy - x^2 y - xy^2 + x^2 y^2$$

$$x(1 - y) + (1 - x)y - x(1 - y)(1 - x)y =$$
$$x + y - 3xy + x^2 y + xy^2 - x^2 y^2$$

Those formulas assume independence of the probabilities involved, and give spurious results when applied to mutually interdependent propositions like these. It happens that none of these formulas gives the best result. If X and Y are independent, the probability that exactly one of them is true is best represented by

$$x + y - 2xy,$$

which is graphed in Figure 8.

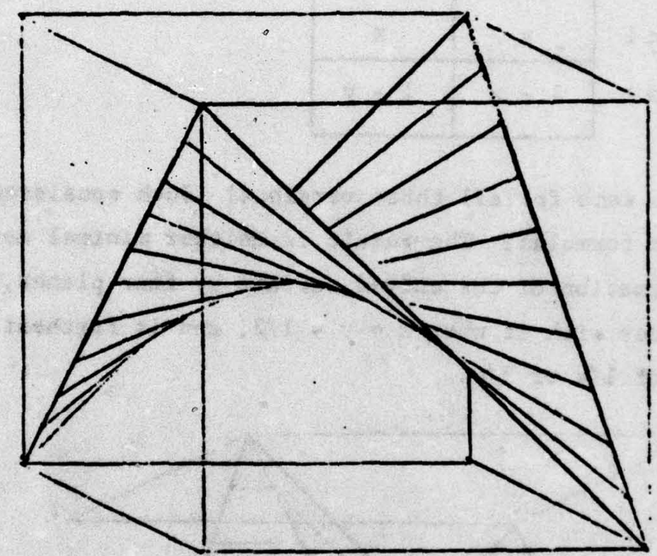

Figure 8. X Exclusive or Y

5.2 Minimax Mixture

Much of what has been written on this subject assumes (sometimes with a brief justification) that the probability of $X \wedge Y$ should always be taken as min (x, y), and that of $X \vee Y$ as max (x, y). What has gone before shows that the former is the maximum possible probability, while the latter is the minimum possible. (It is a source of potential confusion that the one that *is* minimal *uses* a maximum and vice versa.) An interesting thing happens when

these formulas are applied to the various composite representations of "exclusive or":

$$\min(\max(x, y), 1 - \min(x, y))$$
$$\min(\max(x, y), \max(1 - x, 1 - y))$$
$$\max(\min(x, 1 - y), \min(1 - x, y))$$

An easy way to analyze these is to observe that the decisions in them are all governed by whether $x > y$ or by whether $x + y > 1$. So they can be examined for four cases:

	$x < y$	$x > y$
$x + y < 1$	y	x
$x + y > 1$	$1 - x$	$1 - y$

These results are the same for all three versions! Such consistency lends credence to the basic formulas. The result is neither minimal nor maximal, but a pretty fair approximation of the curved surface by four planes, as shown in Figure 9. It coincides with it when $x = y = 1/2$, and is farthest from it when x and y have values of 1/4 or 3/4.

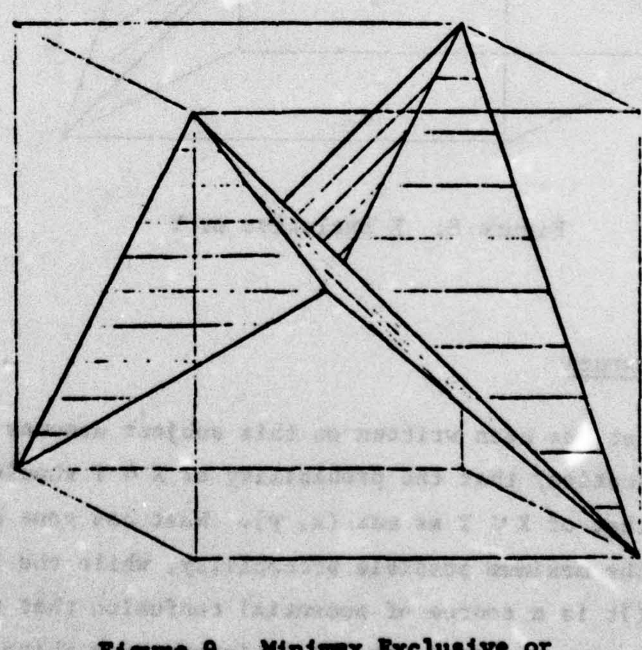

Figure 9. Minimax Exclusive or

5.3 Range

The minimum probability for "exclusive or" is achieved when one of the characteristic sets is included in the other, and is equal to $|x - y|$. The maximum occurs when the sets are as disjoint as possible, and is equal to min $(x + y, 2 - x - y)$. These two values are graphed in Figure 10, enclosing a tetrahedrom that shows a wide range of possibilities. The slices are

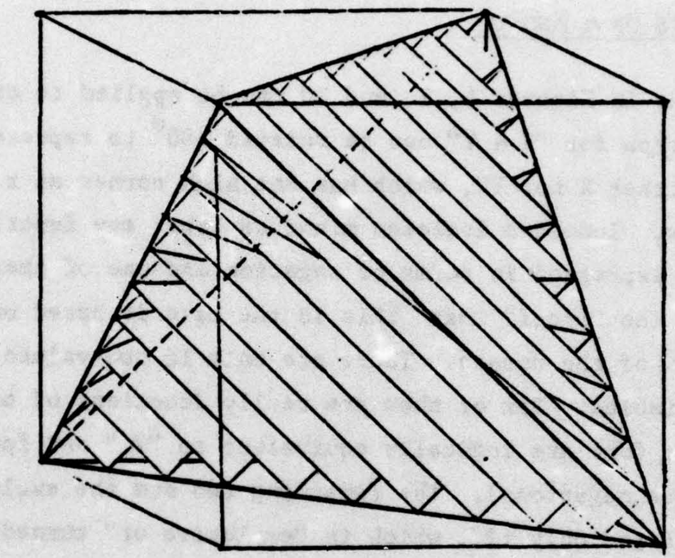

Figure 10. Bounds for Exclusive or

rectangles with their diagonals representing the probable value. The one for $x = .5$ is a square. In the middle of it (where $y = 1/2$ also) the probability s can be anything from 0 to 1, which is as uncertain as a probability can get.

The outer edges of all these graphs are the lines

$x = 0$ and $s = y$

$y = 0$ and $s = x$

$x = 1$ and $s = 1 - y$

$y = 1$ and $s = 1 - x$

They are shared by the maximum, the minimum, the warped surface assuming independence, the four-plane approximation of it, and the two incorrect formulas derived by hasty application of formulas in the text! This shows that all of these formulas are equivalent if <u>either</u> X or Y is restricted to the values 0 or 1. In the propositional calculus, <u>both</u> are so restricted, and so all these formulas are correct for that small subset of the possibilities for two statements.

6. VARIATIONS ON $\wedge$ AND $\vee$

The graphs in Figures 4, 7, and 10 can be applied to other cases. For example, the graph for "X $\wedge$ Y" can be rotated 180° to represent the probability of "neither X nor Y", which has one high corner at $x = y = 0$ and the other three low. Indeed a logician might say that any function of two variables could be expressed in terms of negation and one of these three fundamental functions. The "proof" that this is the case is based on the analysis of only the corners of the domain. There are only 16 two-valued functions of two two-valued variables. Six of them are really functions of only one variable (or none); four are logically equivalent to "$\wedge$" and four more to "$\vee$" (with selective negations). The remaining two are the exclusive or (just considered) and "if and only if", which is "exclusive or" turned sideways.

7. IMPLICATIONS

7.1 $X \rightarrow Y$

These logical equivalences have to do with only the corners of the graphs shown here. Many surfaces can share the same edges, and one of them may be appropriate for different situations from another. An example that deserves particular attention is "if X then Y". Almost every time such a statement shows up in connection with a logical problem, it is summarily replaced by "not X or Y", and then treated like any other "or" statement. This may be just right for a system where everything is two-valued, but it becomes more complicated with continuous variables.

The probability of "$\neg X \vee Y$" is shown in Figure 11, using vertical planes for $x = 0.1, 0.2$, etc. In each plane, a parallelogram shows the maximum probability $\min(1 - x + y, 1)$ and the minimum $\max(1 - x, y)$.

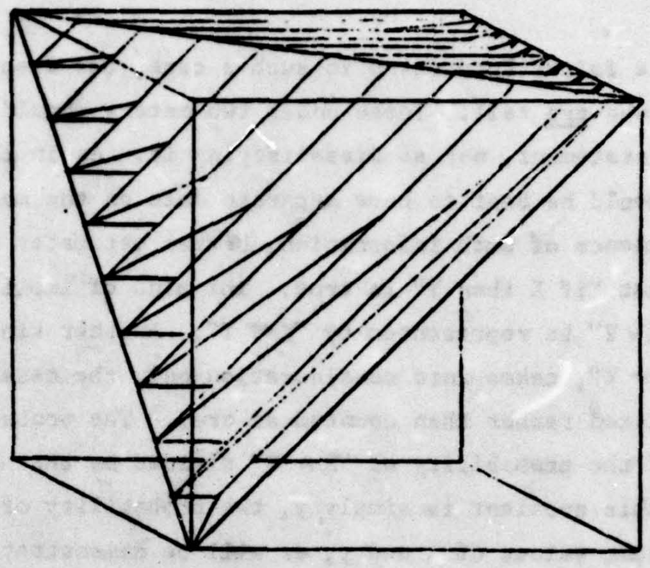

Figure 11. $X \nrightarrow Y$

Between these extremes is the value that assumes independence, $1 - x + xy$, represented by diagonals of the parallelograms, which describe a warped surface as in the other graphs.

7.2 $X \rightarrow Y$

The basic assumption underlying the identification of "if X then Y" with "$\neg X \vee Y$" is that whenever X is false the implication is true by default. Thus either a low probability for X or a high probability for Y contributes to the likelihood that the compound statement is true. If this principle is applied indiscriminately, it can lead to paradoxes. It makes it appear, for example, that this statement is very likely to be true:

> "If a man is over 2 meters tall, his mass is less than 100 kilograms."

Few men are that tall, and most are that light, so that "not tall or light" has a high probability - even higher than "not tall or heavy". But the fact is that the two measures are not independent. The tall men are likely to be the heavy ones.

To make a fairer comparison in such a case, the discussion should be limited to men who <u>are</u> tall. Those under two meters should not be counted as satisfying the statement, nor as dissatisfying it. As in the earlier illustration, it would be best to have separate data on the masses of tall men. But even in the absence of such information, better estimates can be made of the probability that "if X then Y" is true. The kind of implication that is equivalent to "$\neg X \vee Y$" is represented by "$X \Rightarrow Y$". Another kind of implication, represented by "$X \rightarrow Y$", takes into consideration only the cases where X is true. The others are omitted rather than counted as true. The probability of "$X \rightarrow Y$" is the quotient of the probability of "$X \wedge Y$" divided by that of X. If X and Y are independent, this quotient is simply y, the probability of Y. But it can vary greatly for some values of x and y, as will be demonstrated by showing its maximum and minimum.

The probability of "$X \rightarrow Y$" is at least max $(0, 1 - \frac{1-y}{x})$ and at most min $(\frac{y}{x}, 1)$. These limits are graphed in Figure 12. Each plane for a fixed value of x shows a parallelogram reaching from $z = 0$ to 1. The diagonals

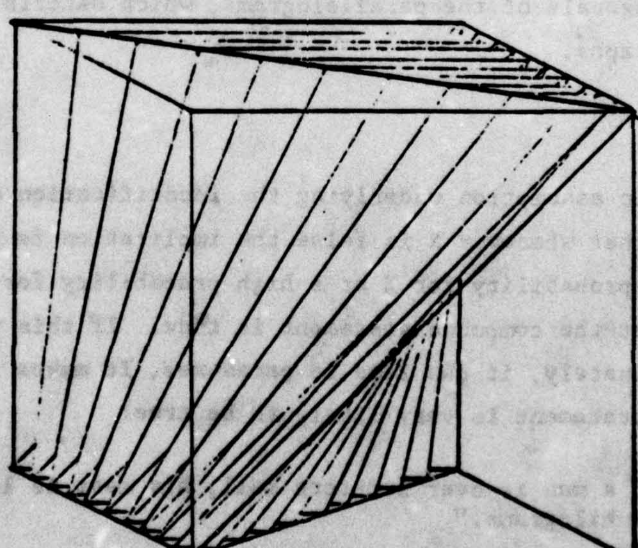

Figure 12. $X \rightarrow Y$

all have the same slope; they form the plane s = y, representing the intermediate value assuming independence of X and Y. For small values of x the parallelograms grow quite large. When x = 0 the graph opens out into a square, showing no restrictions at all on the probability of "X→Y".

7.3 Comparison

Comparison of Figures 11 and 12 shows that "X⇒Y" and "X→Y", while they are identical when x = 1, differ more and more as x approaches 0. One set of limits does not even fit inside the other, as shown in Figure 13.

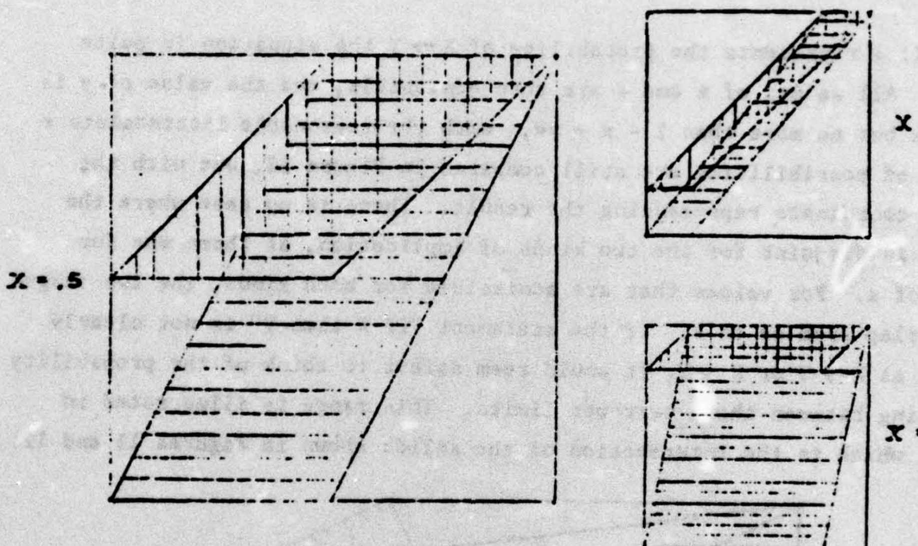

Figure 13. Comparison of X⇒Y and X→Y

Here the slices for x = .5 are superimposed. The limits for "X⇒Y" are cross-hatched vertically, and those for "X→Y" horizontally, and it can be seen that for y < .25 the minimum value for "X⇒Y" is above the maximum for "X→Y". The smaller pictures show how corresponding slices for x = .8 give parallelograms that are close to each other, while those for x = .2 differ greatly.

7.4 Modus Ponens

An important application of the functions graphed in Figures 11 to 13 is the analysis of the probabilities involved in modus ponens. If the statements X and X⇒Y (or X and X→Y) are given, then Y follows as a consequence.

What is the effect of probabilities other than 0 or 1? The answer is already in these graphs. They represent all possibilities regardless of which variable is called dependent. If the probability of X is x, and that of the implication $X \Rightarrow Y$ is e, then Figure 11 shows that $x + e$ must be at least 1. This is because the probability of $X \Rightarrow Y$ (or $\neg X \vee Y$) must be at least as great as that of $\neg X$, that is, $e \geq 1 - x$.

For admissible values of x and e, the probability of Y is at least $x + e - 1$ and at most e, with the reasonable intermediate value $1 - \frac{1 - e}{x} = \frac{x + e - 1}{x}$ derived from the equation of the warped surface $1 - e = x(1 - y)$. This is for an implication $X \Rightarrow Y$.

If e represents the probability of $X \rightarrow Y$ the situation is quite different. All values of x and e are then admissible, and the value of y is at least xe but no more than $1 - x + xe$, with the reasonable intermediate e The ranges of possibilities are still compared in Figure 13, but with the horizontal coordinate representing the result. There is no case where the range of y is disjoint for the two kinds of implication, as there was for the range of e. For values that are admissible for both kinds, the two ranges always overlap from xe to e. If the statement "if X then Y" is not clearly identified as $X \Rightarrow Y$ or $X \rightarrow Y$, it would seem safest to think of the probability of Y as being between these narrower limits. This range is illustrated in Figure 14, which is the intersection of the solids shown in Figures 11 and 12.

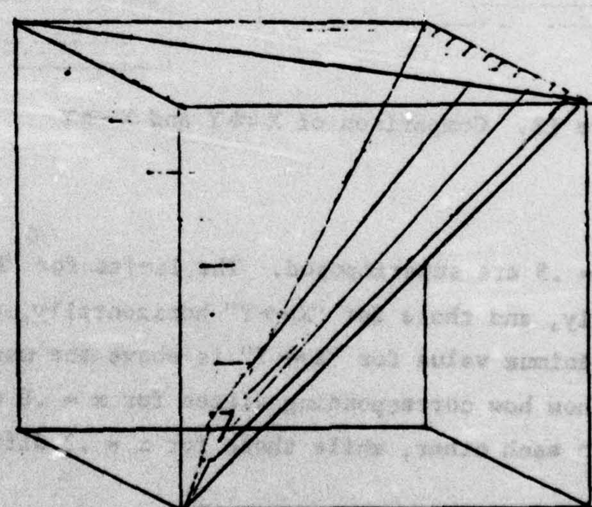

Figure 14. Intersection of $X \rightarrow Y$ and $X \Rightarrow Y$

It has one warped surface (the front), two that are plane triangles, and one triangle with a parabolic notch cut out. Its slices are trapezoids like those doubly cross-hatched in Figure 13.

The comparisons in Figures 11 to 14 are not entirely fair in that they do not show the interdependence of $X \Rightarrow Y$ and $X \rightarrow Y$. The conditions that make s high for *one of* them *also* make it high for the other. They are related by the *formula*

$$(1 - \text{probability of } X \Rightarrow Y) = x \cdot (1 - \text{probability of } X \rightarrow Y)$$

Thus the pairs of parallelograms in Figure 13 are related by vertical stretching or shrinking, with the top edge being held fixed. A particular pair of statements will be represented by points in corresponding positions on the parallelograms. The discrepancy between the values of s is $(1 - x)(1 - s)$, which is largest when x and s are small, but almost negligible when they are both close to 1. A formula that gives satisfactory results for large x and s may be quite inappropriate for small values.

8. RESOLUTION

8.1 The Resolvent

Modus ponens is a special case of deriving *one* statement *from* two others. A more general maneuver is the use *of* the "resolution principle". When a set of information contains the statements:

$$X \vee W \quad \text{and} \quad Y \vee \neg W$$

it is useful to augment the set by the "resolvent" $X \vee Y$. If the individual probabilities of X, Y, and W are known, they determine (or at least narrow the range of) the probabilities of the statements and of the resolvent. A more likely situation is that probabilities will have been assigned to the two compound statements without particular knowledge about their components. If x' is the probability that $X \vee W$ is true, and y' the probability that $Y \vee \neg W$ is true, what is the probability s of $X \vee Y$?

Finding the probabilities x, y, and w from the two numbers x' and y' is tantamount to solving two equations in three unknowns. No definitive answer can be given, even with assumptions of independence. But minima and maxima can be established, and then some observations made about the range of values for the probability of the resolvent.

A key ingredient is w, the probability that W is true. It follows from what has been said before about "or" that

$$w \leq x' \qquad \text{and} \qquad 1 - w \leq y',$$

and these in turn imply that $x' + y' \geq 1$. If the probabilities of the two statements add up to less than 1, it should be concluded that the set of statements is already inconsistent. Such a conclusion is a goal of the resolution procedure, and this shows the beginning of how it works when statements have probabilities other than 0 or 1.

8.2 Range

The bounds of the probability of $X \vee Y$ can be found by considering the probabilities of five statements that are mutually exclusive and exhaustive:

Probability	Statement
p	$W \wedge \neg X \wedge \neg Y$
q	$W \wedge X \wedge \neg Y$
r	$(\neg W \wedge X) \vee (W \wedge Y)$
s	$\neg W \wedge \neg X \wedge Y$
t	$\neg W \wedge \neg X \wedge \neg Y$

The Venn diagram in Figure 15 shows that these five probabilities must add up to 1, and that $p + q + r = x'$ and $r + s + t = y'$. From these facts it follows that $r = x' + y' - 1$. The set with probability r is the intersection of $X \vee W$ with $Y \vee \neg W$. If these were independent statements, the

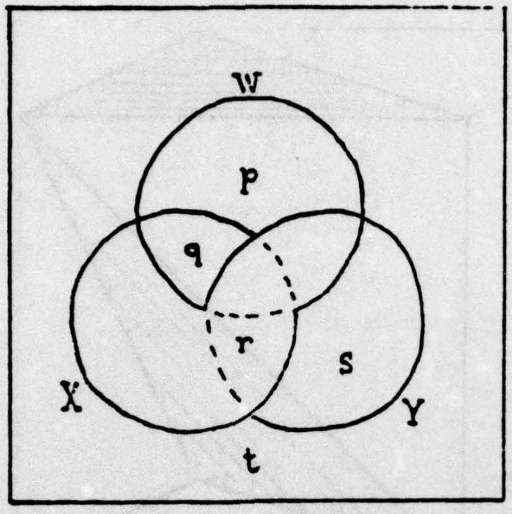

Figure 15, Probabilities in Resolution

probability could range from max $(0, x' + y' - 1)$ to min (x', y'), as demonstrated previously. But the way W is used makes the sets avoid each other as much as possible, so that the minimum is assumed.

The desired probability of $X \vee Y$ is $s = q + r + s$. The restrictions on the probabilities make $p + q = 1 - y'$ and $s + t = 1 - x'$, but do not control the individual values. So it is possible to have a situation in which $p = 1 - y'$, $t = 1 - x'$, and $q = s = 0$, making the probability of $X \vee Y$ as low as possible, namely $s = r = x' + y' - 1$. It can also happen that $p = t = 0$, making s go all the way to 1 for any admissible values of x' and y'. These bounds are graphed in Figure 16, using x' and y' as independent variables rather than x and y.

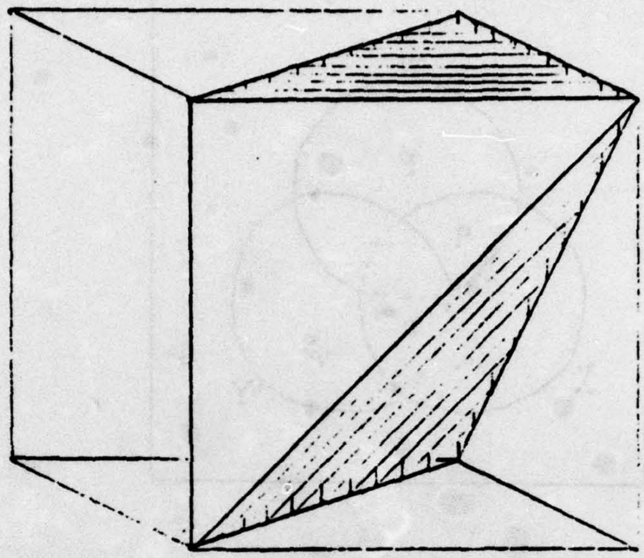

Figure 16. Bounds of Probability of Resolvent

8.3 Intermediate Value

The assignment of a "medium" probability for the resolvent $X \vee Y$ begins with the assumption that W, X and Y are independent. With so many "or" statements, it is easiest to work with the probabilities that they are false. If W and X are independent, the probability that $X \vee W$ is false is $1 - x' = (1 - w)(1 - x)$. The probability that $Y \vee \neg W$ is false is $1 - y' = w(1 - y)$. The probability that the resolvent $X \vee Y$ is false is $(1 - x)(1 - y)$, and the other two equations imply that this is equal to

$$1 - s = \frac{(1 - x')(1 - y')}{w(1 - w)}$$

This has an extra variable w in it, so that it does not yield a single function of x' and y'. But it does incorporate the assumption of independence, so that it may be instructive to see how high or low this variable probability can go. When $w = 1/2$ the denominator of the fraction is largest, and this leads to the highest probability that $X \vee Y$ is true, namely $s = 1 - 4(1 - x')(1 - y')$.

This is graphed in Figure 17. It is a warped surface like previous graphs, but stretched in the vertical direction. It touches $z = 0$ only when $x' = y' = 1/2$.

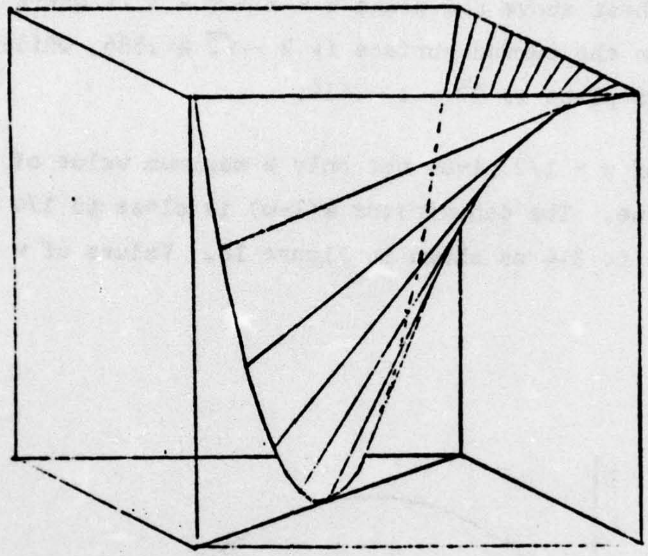

Figure 17. Probability of $X \vee Y$ if $w = 1/2$

It was noted previously that w must be between $1 - y'$ and x'. If both x' and y' are above 1/2, this range will include $w = 1/2$, and that will yield the maximum for z. The minimum will be achieved at one end of the range or the other. If $w = 1 - y'$, $z = \dfrac{x' + y' - 1}{y'}$; if $w = x'$, $z = \dfrac{x' + y' - 1}{x'}$. The lesser of these is the lower bound for z. If either x' or y' is below 1/2, w cannot be 1/2. In those cases the last two formulas represent both ends of the range of z. The maximum for z turns out to be very much like Figure 17. The surface like a pitcher spout pointing to $x' = y' = z = 1$ represents the maximum in the whole quadrant $x' \geq 1/2$, $y' \geq 1/2$. Two warped surfaces tangent to that one have the effect of slightly filling in the parabolic edge where the spout meets the wall $x' + y' = 1$. The spout thus merges into the wall, meeting the floor along the line $x' + y' = 1$.

The minimum for s differs very little from the bottom plane shown in Figure 16. It has the same edges, but is dented slightly where it intersects the plane $x' = y'$. The place where the minimum over all W (assuming independence of X and Y) is furthest above the plane $s = x' + y' - 1$ is where $x' = y' = \sqrt{2}/2 \approx .707$. There the minimum on the curved surface is $2 - \sqrt{2} \approx .586$, while the absolute minimum shown by the plane is $\sqrt{2} - 1 \approx .414$.

The use of $w = 1/2$ gives not only a maximum value of s, but in a sense a typical value. The denominator $w(1-w)$ is close to 1/4 for any value of w from about 1/4 to 3/4 as shown in Figure 18. Values of w closer to the

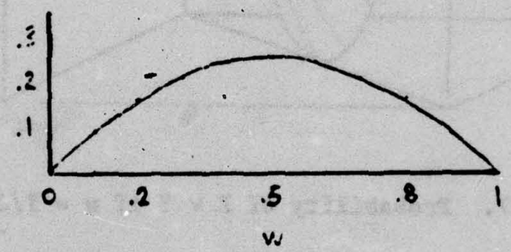

Figure 18. Value of $w(1 - w)$

ends of the scale are likely to be cut off by the limitation that $1 - y' \leq w \leq x'$. So it may be that $s = 1 - 4(1 - x')(1 - y')$ is as good a formula as any for estimating the probability of a resolvent.

9. THE CHOICE

9.1 Minimax

This paper has discussed the full range of possibilities for the elementary functions, and has generally pointed toward the use of the

"independence" formulas: product for $X \wedge Y$, sum-minus-product for $X \vee Y$, etc. These have been seen as a reasonable compromise between the linear minimax functions representing the highest and lowest probabilities. In contrast to this approach is a lot of current literature that assumes the use of $\min(x, y)$ for $X \wedge Y$ and $\max(x, y)$ for $X \vee Y$. The first of these is the maximum possible value, as shown in Figure 2; the second is a minimum, as shown in Figure 7. But the two are compatible with each other, because the probability of $X \vee Y$ is always the sum of the probabilities of X and of Y, <u>minus</u> the probability of $X \wedge Y$. The subtraction makes one go up when the other goes down. Thus another pair of compatible formulas is those at the other end of the range: $\max(0, x + y - 1)$ for $X \wedge Y$, and $\min(x + y, 1)$ for $X \vee Y$. These are not so simple as the other formulas, but technically they deserve as much attention. Why are they not included in the discussions that seem to treat min/max and independence as <u>the</u> two choices?

A possible answer lies in a geometrical view of things that has been illustrated in this paper. Two dimensions represent a domain of interest, and a third is superimposed to show the range of a variable. In these illustrations all three dimensions represent probability, so that it was even legitimate to turn some of the graphs sideways. The domain can have any number of dimensions, but the simplest illustrations result when it has one.

In Figure 19 the horizontal dimension represents a one-dimensional

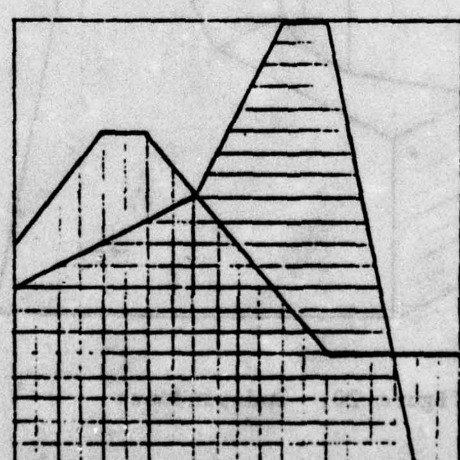

Figure 19. Minimum - Maximum

universe of discourse, and the vertical represents some function defined over that universe. It can be probability or degree of membership in a fuzzy set (which are not the same thing, as L.A. Zadeh pointed out in 1965 - Information and Control, vol 8, page 340). The regions under two rectilinear "curves", cross-hatched in different directions, represent two such functions, which may be called x and y. (It does not matter which is which.) The doubly cross-hatched region shows a third function that is min(x, y). The region that is cross-hatched in either direction (i.e., not white) displays the value of max(x, y). Variations on this figure appear often in discussions of fuzzy sets, and are used in defense of the adoption of minimum for "and" (or intersection) and maximum for "or" (or union).

9.2 Independence

The "other" formulas, xy and x + y - xy, can be visualized by a mathematician looking at Figure 19, but do not stand out the way min and max do. But independence can be better represented by using two different dimensions for x and y. In Figure 20 one of them is shown vertically as before,

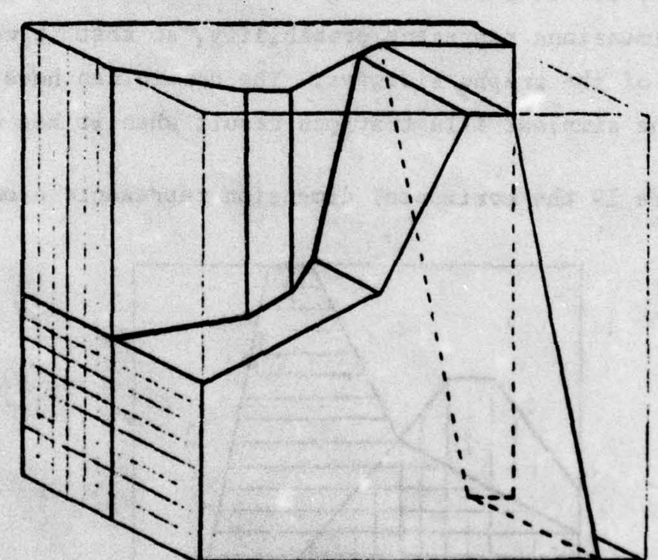

Figure 20, Independence

but the other is plotted from back to front. For every element in the universe there is a plane section of this cube, typified by the left end that shows in Figure 20, where the two probabilities happen to be .4 and .5. The

doubly cross-hatched region has an area of .2 of the whole square, or in general of the product xy. The L-shaped non-white region represents x + y - xy (.7 at the left end). The intersection of the two surfaces shows where the corner of the doubly cross-hatched rectangle is for all elements in the universe. In this illustration the minimum and maximum can be visualized with some effort, but they are upstaged by the independence formulas.

9.3 The Other Extreme

Much of the discussion of this subject treats these two as if they were the only choices. As has been shown in this paper, there is really a whole spectrum of possibilities. Min and max are at one end of it, and independence in the middle. Why has the discussion not included the other end of the spectrum? It may be because it has not been given an appealing geometrical representation. But there is one available.

In Figure 21 are shown the same functions as in the previous two figures. They are both plotted as vertical coordinates, but they start from

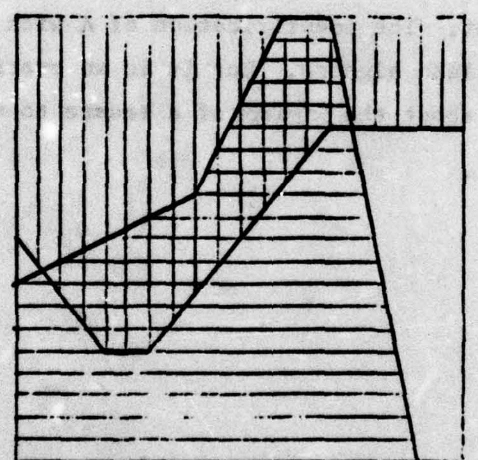

Figure 21. The Other Extreme

opposite ends of the scale! In this unfamiliar position the hitherto neglected formulas have their moment of glory. The height of the doubly cross-hatched

region is max(0, x + y - 1), the lowest possible value for $X \wedge Y$. The combined height of the non-white regions is min(x + y, 1), which is the highest possible value for $X \vee Y$.

It may well be asked, But doesn't using opposite ends of the scale introduce a strong bias? The answer is, Yes, and so does using the same end of the scale! As natural as it may seem for both graphs to be "upright," it makes X and Y overlap as much as they possibly can. Thus two statements with probabilities ρf .1 each are displayed as coinciding in spite of all the room they had to miss each other! Similarly, the opposite-end style makes two .5's miss each other completely, which is pretty much of a fluke too. Both representations are extreme because they are showing the largest and smallest possibilities.

An appealing thing about such diagrams as these is that they present perceptible geometric sets to represent abstractions. The unions and intersections of those sets seem very naturally to represent corresponding functions such as disjunctions or conjunctions of statements. The intersection in Figure 14 represents a yielding to the temptation to work with the graphs as entities in their own right. The identification of $\wedge$ with $\cap$ and $\cup$ with $\vee$ seems like good, clean Boolean algebra. But it is an oversimplification, comparable to using facts about the corners of a square to derive conclusions about its whole interior.

10. A CREDIBILITY FORMULA FOR COMPOUND IMPLICATIONS

If $A \wedge B \wedge C \wedge \ldots \Rightarrow Q$, and if all the antecedents A, B, C, etc., are true, then Q follows as a consequence. The antecedents, and possibly the implication itself, may have associated with them numbers in the interval $[0,1]$ representing degrees of credibility. These can be thought of as probabilities, but other interpretations are also possible. This paper presents a way of assigning a credibility number to Q, depending on the credibilities of the other statements: $\langle Q \rangle = f(\langle A \rangle, \langle B \rangle, \langle C \rangle, \ldots \langle A \wedge B \wedge C \wedge \ldots \Rightarrow Q \rangle)$.

If $\neg A$ is the contradiction of A, the simplest interpretation dictates that $\langle \neg A \rangle = 1 - \langle A \rangle$. But this is not the only possibility. In dealing with uncertain sources of information, it may be useful to think of $\langle A \rangle$ and $\langle \neg A \rangle$ as two numbers whose sum is less than 1, with the defect representing a range of uncertainty. This possibility will not be explored in this note, but $\langle A \rangle$ and $\langle \neg A \rangle$ will be represented separately, without any assumption that either one determines the other.

The complexity of the problem to be considered here is evident from the "simplest" case, in which there is only one antecedent. The credibilities of $\langle A \rangle$ and $\langle A \Rightarrow Q \rangle$ do not very closely control that of $\langle Q \rangle$. For example, $\langle A \Rightarrow Q \rangle$ might represent the statement, "If a number is prime, it is odd." This statement has a credibility slightly less than 1, because the number 2 is prime but even. Suppose it has been determined that the probability that N is prime is .6; what is the probability that N is odd?

If N really is prime it is very probably odd, and this suggests a probability just under .6; but even if N is not prime it has a chance of being odd, and this may add almost .2 more to the chances, for a total just under .8.

In any situation like this, it would be useful to know two credibilities, which are best described in terms of probabilities, but can be applied also to other interpretations of credibility. The universe of discourse contains four

kinds of elements, with their total numbers represented by the letters in this chart. The total $w + x + y + z$ is taken to be 1.

	Q true	Q false
A true	w	x
A false	y	z

The probability of $A \Rightarrow Q$ is $w + y + z$. That of A is $w + x$. Knowing these two numbers does not make it possible to find $\langle Q \rangle = w + y$. But the conditional probability of $A \rightarrow Q$ is $\frac{w}{w + x}$, and that of $\neg A \rightarrow Q$ is $\frac{y}{y + z}$. If both of these numbers are known, the probability of the consequence is given explicitly by the formula

$$\langle Q \rangle = \langle A \rangle \langle A \rightarrow Q \rangle + \langle \neg A \rangle \langle \neg A \rightarrow Q \rangle.$$

If a rule of the form $A \Rightarrow Q$ can have two such numbers supplied with it, good estimates can be made of the credibilities of consequences.

The A in the above discussion could be taken to represent a conjunction of several antecedents, so that the whole discussion would apply to other situations. But another illustration can show that this might be hasty. Consider the rule:

$$(x \text{ works with } y) \wedge (y \text{ works at } z) \Rightarrow (x \text{ works at } z).$$

A single credibility for this rule would be quite high, falling short of 1 only because of such anomalies as people who work on more than one thing. The other credibility that ought to be known is that of the supplementary statement:

$$(x \text{ does not work with } y) \vee (y \text{ does not work at } z) \rightarrow (x \text{ works at } z \text{ anyhow}).$$

How can one number describe such a probability? If x does work with y but y does not work at z, there is practically no chance that x works at z. If, on the other hand, x does not work with y, while y does work at z, it still might

very well be true that x works at s. And if both antecedents are false, the probability could be almost anything, depending on the universe to which x belongs. Now, if <u>four</u> probabilities could be given in connection with such a rule, the problem would be solved:

$$
\begin{aligned}
\langle Q \rangle = {} & \langle A \wedge B \rangle \langle A \wedge B \rightarrow Q \rangle + \\
& \langle A \wedge \neg B \rangle \langle A \wedge \neg B \rightarrow Q \rangle + \\
& \langle \neg A \wedge B \rangle \langle \neg A \wedge B \rightarrow Q \rangle + \\
& \langle \neg A \wedge \neg B \rangle \langle \neg A \wedge \neg B \rightarrow Q \rangle .
\end{aligned}
$$

For three antecedents there would be eight terms in such an expression, for four antecedents sixteen, and for n antecedents 2^n. This very quickly gets cumbersome, and will not be pursued. But one thing will be salvaged before it is dropped.

The probability of $\langle A \wedge B \rangle$ is not simple itself, involving a conditional probability. But it has been demonstrated in 2330-TN-1 that the assumption of independence simplifies this without being very likely to do serious damage. The assumption that $\langle A \wedge B \rangle = \langle A \rangle \langle B \rangle$ makes the ensuing discussion very much simpler than it would be without it, and will accordingly be adopted temporarily. Later it will be discussed how to cope with interdependence of antecedents.

The case of three antecedents ($A \wedge B \wedge C \rightarrow Q$) will be used as a vehicle to discuss a technique that applies to any number (even 1 or 2).

The simplest approach is to supply one credibility for the rule:

$$Z_o = \langle A \wedge B \wedge C \rightarrow Q \rangle$$

For some rules Z_o may be 1, but it can be considerably lower, as in the exhortation: "If you stand right and keep your eye on the ball and swing level you will hit it."

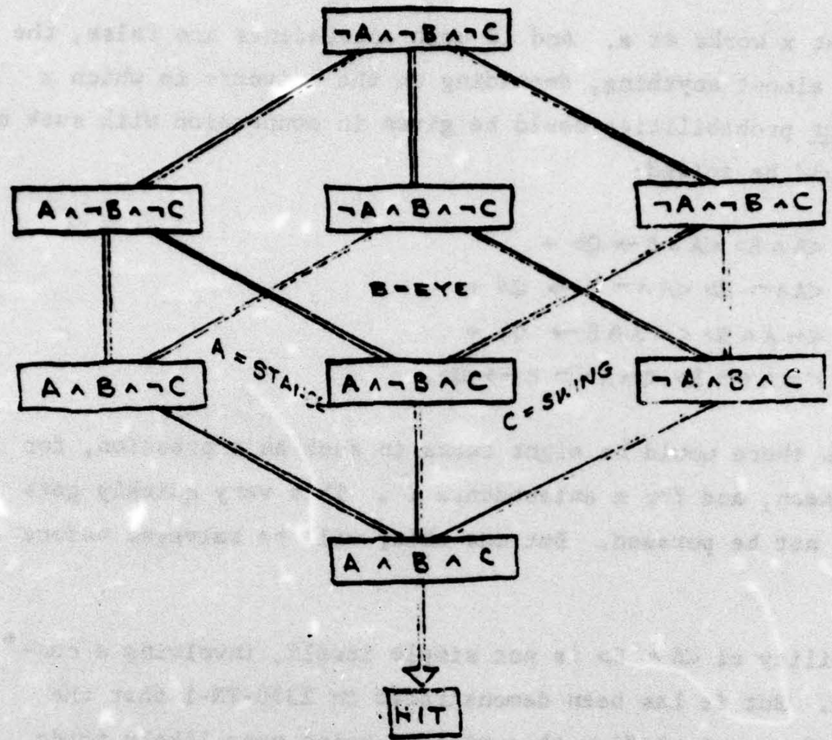

The meeting of the three conditions corresponds to being at one vertex of a cube (or in general an n-dimensional orthotope). Seven other vertices (or $2^n - 1$) represent various kinds of poor technique, the worst of which is the opposite vertex. But even a batter with his feet in the bucket <u>and</u> a bad eye <u>and</u> an awkward chop has <u>some</u> chance to hit the ball. (Whether he will pop up, ground out, etc., is another question.) The coach may be able to supply eight probabilities, each to be multiplied by the probability of its combination, such as <A> <¬B> <C> for good stance, bad eye, good swing (multiplication being used on the assumption of independence). It turns out that these 2^n numbers can be replaced by only n + 1 (including Z_o).

What is called for is an estimate of the results of each single kind of failure. In the example, these might be:

$Z_A = .6$ poor stance reduces hit probability to .6 of what it would have been with good stance;

$Z_B = .3$ eye off the ball reduces to .3;

$Z_C = .5$ improper swing cuts chances in half.

The new Z's are not probabilities all by themselves, but factors that reduce the overall probability Z_o:

$$Z_A\ Z_o = \langle \neg A \wedge B \wedge C \rightarrow Q \rangle$$

$$Z_B\ Z_o = \langle A \wedge \neg B \wedge C \rightarrow Q \rangle$$

$$Z_C\ Z_o = \langle A \wedge B \wedge \neg C \rightarrow Q \rangle.$$

If, for example $Z_o = .8$, the three mediocre batters have hit probabilities of .48, .24, and .40, respectively.

The probabilities for the three poor batters who do two things wrong, and the one really bad one, are taken care of automatically by this formula:

$$\langle Q \rangle = Z_o\ (\langle A \rangle + Z_A \langle \neg A \rangle)(\langle B \rangle + Z_B \langle \neg B \rangle)(\langle C \rangle + Z_C \langle \neg C \rangle).$$

Each Z_x is multiplied by the probability of failure $\langle \neg X \rangle$, and the product increased by the probability of success $\langle X \rangle$. These sums are all multiplied together along with Z_o.

This is the working version of the formula, involving 2n multiplication and n additions. The expansion below is not intended to show an alternative procedure, but to display the result in a way that clarifies how the formula parlays n + 1 credibilities into 2^n. The algebraic multiplication gives this product:

$$\begin{aligned}\langle Q \rangle = Z_o\ (&\langle A \rangle \langle B \rangle \langle C \rangle + Z_A \langle \neg A \rangle \langle B \rangle \langle C \rangle + Z_B \langle A \rangle \langle \neg B \rangle \langle C \rangle + \\ &Z_C \langle A \rangle \langle B \rangle \langle \neg C \rangle + Z_A\ Z_B \langle \neg A \rangle \langle \neg B \rangle \langle C \rangle + Z_A\ Z_C \langle \neg A \rangle \langle B \rangle \langle \neg C \rangle + \\ &Z_B\ Z_C \langle A \rangle \langle \neg B \rangle \langle \neg C \rangle + Z_A\ Z_B\ Z_C \langle \neg A \rangle \langle \neg B \rangle \langle \neg C \rangle)\end{aligned}$$

The factors $\langle A \rangle$, $\langle B \rangle$, $\langle C \rangle$ appear with every combination of negations, making a full 2^n terms. If each $\langle \neg X \rangle = 1 - \langle X \rangle$, these products all add up to 1. Each one is multiplied by the product of the Z's for all the failures it involves. This carries through the idea of independence giving good medium sized estimates for probabilities. The intuitive appeal of the results can be seen in the batting illustration, where the coefficients are:

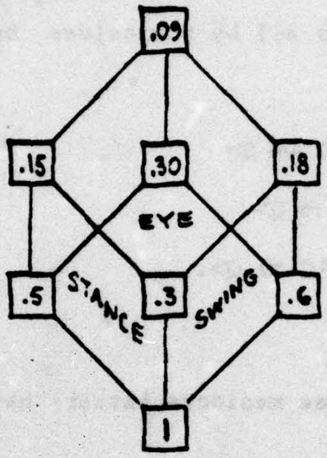

Each one is multiplied by the overall $Z_o = .8$, giving the worst batter only .072 of a chance of putting the wood on it, while even the best gets only .8 of a chance.

A "really good" rule has a high value of Z_o, signifying that if the conditions are met the consequence is really likely to come off. In some cases, like "x works with y $\rightarrow$ y works with x" it is appropriate that $Z_o = 1$. Another characteristic of a "good" rule is low values of the other Z_x numbers. If a Z_x is high, its corresponding condition is not really important. The batter might think, for example, that he should have a rabbit's foot in his pocket (D), but objective analysis might reveal that $Z_D = 1$. On the other hand, having a bat in his hand (E) is <u>so</u> important that it could be taken for granted, and $Z_E = 0$.

Similarly, "really good" data has high values of $\langle X \rangle$ and low values of $\langle \neg X \rangle$. This means that in a system with good information, most of the terms of the above expansion will be microscopic. It is tempting to use this as an excuse for an oversimplified formula. The one given here appears to give the

"negligible" quantities their due, and this is likely to yield benefits in any system. But it becomes especially valuable when the rules and data are "not so good".

The formula works for all values of n. The number of operations involved in it grows linearly, but the effective number of terms (as shown in the expanded version) grows exponentially. It should be noted that for $n = 1$, the interpretation is slightly different from what was said above about that case. Z_o is the same as $\langle A \rightarrow Q \rangle$. The supplementary probability $\langle \neg A \rightarrow Q \rangle$ is not Z_A, but $Z_A Z_o$. In general, $\langle \neg A \rightarrow Q \rangle$ will be low for a good rule, so that dividing $\langle \neg A \rightarrow Q \rangle$ by Z_o will give a value for Z_A that is below 1.

The formula is commutative in the sense that the conditions may appear in any order. The different Z's will do their job without any restriction such as treating them in order of their "goodness". There is also an important sense in which the formula is associative.

A conclusion may often appear as a consequence of many conditions, even in a simple system where the rules never have more than two. If an analysis uses the statements $A \wedge B \Rightarrow P$ and $P \wedge C \Rightarrow Q$ to arrive at the conclusion Q, it may well be said that Q is true because A and B and C are true. It is interesting to see what distinction appears in the use of this formula.

Suppose that the rule $A \wedge B \Rightarrow P$ has credibility numbers Z_o, Z_A, and Z_B, as defined above, and the rule $P \wedge C \Rightarrow Q$ uses Z_1 (in place of Z_o), Z_p, and Z_C. Then:

$$\langle P \rangle = Z_o\,(\langle A \rangle + Z_A \langle \neg A \rangle)(\langle B \rangle + Z_B \langle \neg B \rangle);$$

$$\langle Q \rangle = Z_1\,(\langle P \rangle + Z_p \langle \neg P \rangle)(\langle C \rangle + Z_C \langle \neg C \rangle) =$$

$$Z_1\,(Z_o(\langle A \rangle + Z_A \langle \neg A \rangle)(\langle B \rangle + Z_B \langle \neg B \rangle) + Z_p \langle \neg P \rangle)(\langle C \rangle + Z_C \langle \neg C \rangle).$$

Now, if the second rule is "really good" with respect to P, Z_p will be small, possibly even zero. If $Z_p = 0$, the expression for $\langle Q \rangle$ simplifies to just what it would have been for one rule $A \wedge B \wedge C \Rightarrow Q$, using the same values for Z_A, Z_B, and Z_C, but with the overall credibility $Z_o Z_1$, the product of those for the

two individual rules. A non-zero value for Z_P will raise this credibility, reflecting the fact that Q follows not only from $A \wedge B \wedge C$, but from an additional domain where P is true even though A and B are not both true.

This way of combining two rules into one can also be applied to separating one rule into two. It may be useful to replace

$$A \wedge B \wedge C \wedge D \wedge E \rightarrow Q$$

by two rules: $D \wedge E \rightarrow K$ and $A \wedge B \wedge C \wedge K \rightarrow Q$. One place this may have practical value is in the case that D and E are decidedly <u>not</u> independent, so that they upset the validity of the overall scheme. If they are pulled out of the rest of the set, they can be analyzed separately:

$$\begin{aligned}\langle K\rangle = \;& \langle D \wedge E\rangle \langle D \wedge E \rightarrow K\rangle \; + \\ & \langle D \wedge \neg E\rangle \langle D \wedge \neg E \rightarrow K\rangle \; + \\ & \langle \neg D \wedge E\rangle \langle \neg D \wedge E \rightarrow K\rangle \; + \\ & \langle \neg D \wedge \neg E\rangle \langle \neg D \wedge \neg E \rightarrow K\rangle\end{aligned}$$

The atomic probabilities $\langle D \wedge E\rangle$, etc., can be found other ways than by the product. For example, if $\langle D\rangle$ and $\langle E\rangle$ are both high, and the two statements D and E are decidedly negatively correlated, it might be appropriate to use:

$$\begin{aligned}&\langle D \wedge E\rangle = \langle D\rangle + \langle E\rangle - 1\\ &\langle D \wedge \neg E\rangle + 1 - \langle E\rangle\\ &\langle \neg D \wedge E\rangle = 1 - \langle D\rangle\\ &\langle \neg D \wedge \neg E\rangle = 0\end{aligned}$$

The associated conditional probabilities could be supplied for these two variables out of context far more easily than when they are tangled up with A, B, and C also. Then making $Z_K = 0$ would make this pair of rules just like the original single rule. But some other value of Z_K might be chosen for even better accuracy.

APPENDIX C

CONDITIONAL PROBABILITY IN RULES

CONDITIONAL PROBABILITY IN RULES

The most practical method of obtaining the credibility $\langle r \rangle$ of a derived fact r from the rule:

$$p \wedge q \Rightarrow r \qquad \text{given } \langle p \rangle \text{ and } \langle q \rangle$$

consists of treating the hypothesis $p \wedge q$ as a single fact hypothesis. Thus:

$$\langle r \rangle = \langle (p \wedge q) \rightarrow r \rangle \cdot \langle p \wedge q \rangle + \langle \neg (p \wedge q) \rightarrow r \rangle \cdot \langle \neg (p \wedge q) \rangle$$

The two conditional probabilities are analogous to the two inference rule credibility numbers and are expected to appear normally in the STIS information. Thus:

$$(\langle \neg (p \wedge q) \rightarrow r \rangle, \ \langle (p \wedge q) \rightarrow r \rangle) \qquad p \wedge q \Rightarrow r$$

Thus, the $\langle r \rangle$ computation is simple, most particularly in the case where $\langle p \wedge q \rangle = \langle p \rangle \cdot \langle q \rangle$ because of independence or the likely case that no conditional probability is known linking p with q. Here we are concerned with the less usual situation where more complete conditional probability information is available. We adopt the notation:

$$(m, n) \qquad p \Rightarrow q$$

to mean

$$\langle \neg p \rightarrow q \rangle = m \quad \text{and} \quad \langle p \rightarrow q \rangle = n$$

Consider the example:

$$(.15, .9) \qquad p \Rightarrow q \qquad \qquad \langle p \rangle = .85$$

$$p \wedge q \Rightarrow r$$

and the task is to estimate $\langle r \rangle$. As an intermediate step we estimate $\langle p \wedge q \rangle$ which is always given by the conditional probability law:

$$\langle p \wedge q \rangle = \langle p \rangle \ \langle p \rightarrow q \rangle$$

We observe, now, that the .9 number is the conditonal probability

$$\langle p \rightarrow q \rangle$$

Therefore, our work is simply:

$$\langle p \wedge q \rangle = \langle p \rangle \langle p \rightarrow q \rangle = (.85)(.9) = .765$$

Since we have assumed our main rule ($p \wedge q \Rightarrow r$) as being of full certainty; we have:

$$p \wedge q \rightarrow r \quad \text{or} \quad (0, 1)\ p \rightarrow q$$

which is to say that the consequent never fails when the hypothesis is true. We also have (as presented above) the consequent necessarily false when the hypothesis is false. Therefore, we obtain:

$$\langle r \rangle = (1)(.765) = .765$$

Now we proceed with the second matter, namely the credibility features intrinsic to the main implication itself:

$$p \wedge q \rightarrow r$$

We develop the previous pattern for credibility numbers in a rule using as an example:

$$\begin{aligned} \text{where } \langle r \rangle &= .95 \quad \text{when} \quad \langle p \rangle = \langle q \rangle = 1, \text{ or } .95 = \langle p \wedge q \rightarrow r \rangle \\ \langle r \rangle &= .1 \quad \text{when} \quad \langle \neg p \rangle = \langle q \rangle = 1, \text{ or } .1 = \langle \neg p \wedge q \rightarrow r \rangle \\ \langle r \rangle &= .05 \quad \text{when} \quad \langle p \rangle = \langle \neg q \rangle = 1, \text{ or } .05 = \langle p \wedge \neg q \rightarrow r \rangle \end{aligned}$$

We make an initial estimate for $\langle r \rangle$ based on the truth estimate of the hypothesis alone.

$$\text{initial estimate } \langle r \rangle = (.95)\langle p \wedge q \rangle = (.95)(.765) = .727$$

It is the hypothesis being false condition that is more difficult to account for.

We notice that our conventions imply that the rule:

$$(.15,\ .9) \qquad p \rightarrow q$$

has constants to be interpreted thus:

$$(\langle \neg p \rightarrow q \rangle,\ \langle p \rightarrow q \rangle) \qquad p \rightarrow q$$

which suggests that our $\langle r \rangle$ estimates can be arrived at as the following sum:

$$\begin{aligned}
&\text{initial estimate} && \langle r \rangle = (.95)\ \langle p \wedge q \rangle && = (.95)(.765) && = .727 \\
&\text{first correction for} && \langle r \rangle = (.1)\ \langle \neg p \rightarrow q \rangle \langle \neg p \rangle && = (.1)(.15)(.15) && = .002 \\
&\text{second correction for} && \langle r \rangle = (.05)\ \langle p \rightarrow \neg q \rangle \langle p \rangle && = (.05)(1-.9)(.85) && = \underline{.004} \\
&\text{Final estimate for} && \langle r \rangle && && = .733
\end{aligned}$$

We note that in the second correction is a result of noting that we start with the condition $\langle p \rangle = \langle \neg q \rangle = 1$ for which the probability $\langle p \rightarrow \neg q \rangle$ is appropriate. Then we notice that $\langle p \rightarrow q \rangle$ comes from the simpler implication, and that $\langle p \rightarrow q \rangle + \langle p \rightarrow \neg q \rangle = 1$.

As with the single literal hypothesis case, a Venn diagram is given with shadings indicating the two different corrections applied to the original $\langle r \rangle$ estimate. The unshaded portion inside r represents the initial estimate for $\langle r \rangle$ (.727).

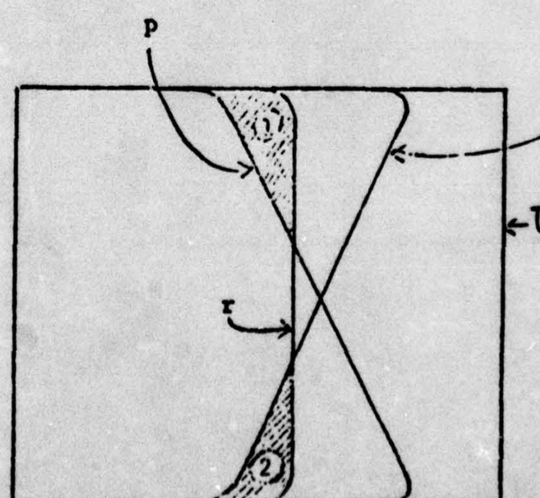

We also note that every section of this Venn diagram has an appropriate interpretation in the above work estimating $\langle r \rangle = .735$, which is diagrammatically underscaled. We list a few, as follows:

$\langle p \rangle = .85$	assumption, also underscaled
$\langle p \rightarrow q \rangle = .9$	assumption, underscaled
$\langle \neg p \rightarrow q \rangle = .15$	assumption, overscaled
$\langle p \wedge q \rangle = .765$	calculated, underscaled
initial $\langle r \rangle = .727$	calculated, unshaded, underscaled
first correction $\langle r \rangle = .0C2$	calculated, marked 1, overscaled
second correction $\langle r \rangle = .004$	calculated, marked 2, overscaled

We further note that a third possible correction for $\langle r \rangle$, in the event that both p and q individually fail, has been neglected.

APPENDIX D

ILLUSTRATIONS OF THE USE OF REPORT LIKELIHOOD

ILLUSTRATIONS OF THE USE OF REPORT LIKELIHOOD

VENN DIAGRAMS

We suppose the fact:

.75, WA(2,7)

might appear diagrammatically thus:

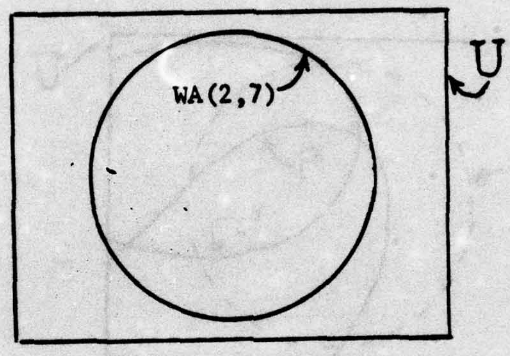

where the fact WA(2,7)) appears as a circle surrounding a truth area comprising 75% of the universe area. Now we proceed to apply this graphical scheme to the credibility ideas pertaining to system reports.

In our illustration we develop the likelihood ratios of two reports on the fact WA(2,7), namely:

$$L_{R_1} = \frac{\langle WA \rightarrow R_1 \rangle}{\langle WA \rightarrow R_1 \rangle} = \frac{.75}{1 - .75} = 3$$

$$L_{R_2} = \frac{\langle WA \rightarrow R_2 \rangle}{\langle WA \rightarrow R_2 \rangle} = \frac{.80}{1 - .80} = 4$$

These are based on the two reports introduced as follows:

$$.75,\ R_1 \rightarrow WA(2,7)$$

$$.80,\ R_2 \rightarrow WA(2,7)$$

where the credibility numbers are the usual probabilities for the truth of the WA fact based on the appropriate report alone (under the important assumption that <WA(2,7)> starts from an old credibility estimate of .50, prior to both reports.)

We represent this situation with the following Venn diagram:

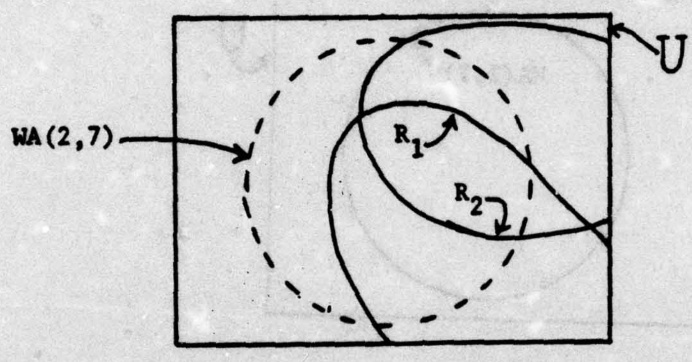

in which the $\langle \overline{WA} \rightarrow R_1 \rangle$ probability is represented by the enclosure of 1/4 the $\overline{WA}$ area inside the R_1 loop. The $\langle WA \rightarrow R_1 \rangle$ probability is represented by the enclosure of 3/4 the WA area inside the R_1 loop. Thus, the likelihood ratio indicates the degree to which the R_1 loop prefers the truth value set of the fact WA(2,7). When $L[WA(2,7) \rightarrow R_1] = 1$, the R_1 report may be though of as independent of the WA fact, that is, the R_1 report shows no preference between the two sides of the WA(2,7) oval. When $L[WA(2,7) \rightarrow R_1]$ is large, we have the R_1 report nearly inside the WA truth set. This is the idealistic situation where the report R_1 is highly reliable, and the Bayes result is:

$$L_{new}[WA(2,7)] = L[WA(2,7) \rightarrow R_1]\ L_{old}[WA(2,7)]$$
$$= L[WA(2,7) \rightarrow R_1]\ (1)$$

meaning that <WA(2,7)> is almost unity and the fact is assured a high likelihood.

We notice that the same state of affairs can also be represented by a new R_1 oval which encloses only 1/4 the WA area and 3/4 the $\overline{WA}$ area. This may be thought of as having the same value and the same likelihood ratio as the earlier R_1 report. The difference between the two may be thought of resting in how appropriate the report is to appear and in that alone.

When the two loops are nearly identical (reports duplicate substantially), then the intersection area $R_1 \wedge R_2$ gives no substantial likelihood ratio improvement (the preference for the WA truth set is relatively unchanged). The Venn diagram as actually drawn indicates high independence of reports R_1 and R_2 with a joint likelihood ratio approximately:

$$L_{R_1 \wedge R_2} = (L_{R_1})(L_{R_2})$$
$$= 3 \times 4 = 12$$

Such a computation is exactly the result of applying the Bayes technique consecutively, with the normal assumption of report independence.

INITIAL FACT LIKELIHOOD

In the text of this report we have made a simplifying assumption that the historically established credibility estimate is:

$$\langle WA(2,7)\rangle_{old} = .5$$

before the appearance of the two reports R_1 and R_2 are to be accounted for. The simplification is a result of the use of the resulting L_{old} [WA(2,7)] = 1 in the Bayes relation:

$$L_{new}\left[WA\right] = L\left[WA(2,7) \rightarrow R_1\right] L_{old}\left[WA\right] \qquad (R_1 \text{ only})$$

$$L_{new}\left[WA\right] = \frac{\langle WA(2,7)\rangle_{new}}{\langle \overline{WA}(2,7)\rangle_{new}} = L\left[WA(2,7) \rightarrow R_1\right] = \frac{\langle WA \rightarrow R_1\rangle}{\langle \overline{WA} \rightarrow R_1\rangle} \qquad (R_1 \text{ only})$$

in which the R_1 report can be characterized in each of two equivalent ways:

$$L\left[WA(2,7) \rightarrow R_1\right] = 3 \; (= \frac{.75}{1 - .75})$$

$$\therefore \; \langle R_1 \rightarrow WA(2,7)\rangle = .75$$

The simplification is only computational, as we will illustrate below.

We do the same problem with the following alteration:

$$\langle WA(2,7)\rangle_{old} = .05, \text{ instead of } .5$$

If facility #7 is one out of a total of 20 facilities, the new assumption represents an approximate zero information state. The previously assumed value of .5 really represents a lot of information: person #2 is as likely (or more so) to work at facility #7 as at any one amongst all the remaining 19 together. Let us see what is the result of applying the same R_1, R_2 reports with this new fact assumption. We get from two applications of the Bayes rule:

$$L_{new}\left[WA\right] = L\left[WA \rightarrow R_2\right] \; L\left[WA \rightarrow R_1\right] \; L_{old}\left[WA\right]$$

$$L_{new}\left[WA\right] = (8)(4)(\frac{.05}{.95}) = \frac{12}{19}$$

This means that, after the R_1, R_2 reports are accounted for, we have:

$$\frac{\langle WA(2,7)\rangle_{new}}{1 - \langle WA(2,7)\rangle_{new}} = \frac{12}{19}$$

$$\langle WA(2,7)\rangle_{new} = \frac{12}{19 + 12} = \frac{12}{31}$$

We notice that the WA credibility still falls short of the 1/2 level. One more report of value approximately that of R_1 or R_2 will suffice. We can see what sort of R_3 would be required by reapplying the Bayes Theorem in a different manner, where it is now:

$$L_{old}\left[WA\right] = \frac{12}{19} \text{ (not } L_{new}\text{)}:$$

$$L_{new}\left[WA\right] = 1 = L\left[WA \rightarrow R_3\right] \; (\frac{12}{19})$$

Therefore: $L\left[WA \rightarrow R_3\right] = \frac{19}{12}$

Since $\frac{19}{12} < 3 < 4$

We see that R_3 is less than R_1 or R_2 in value, to bring the final fact credibility to just 1/2.

INFORMATION MEASURE

Since we have already identified the value of a report with its likelihood ratio, we already have an indicator of the amount of information of a report. The more valuable a report, the more information. However, we do have a situation where two reports with likelihood ratios of 3 and 4 have the same information content (or value) as one report with likelihood ratio 12. Another scale is possible, much as is used in psychological measurements.

We consider the information measure I_R for a report R as defined to be:

$$I_R = \text{Ln} \quad (L\left[WA \rightarrow R\right])$$

We tabulate the following results:

Report	$L\left[WA \rightarrow R\right]$	$I_R = \text{Ln} \quad (L\left[WA \rightarrow R\right])$
R_1	3	1.10
R_2	4	1.39
useless report	1	0
(R_1 and R_2)	12	2.49

in which we see that the information measures add up pleasantly. One unit of information in a report means that the likelihood ratio of a system fact is increased by a factor of $e \doteq 2.718$.

This emphasizes that an individual report is well identified by the amount of information it brings, or alternatively by the report likelihood ratio. For example, our first R_1 report was characterized thus:

$$L\left[WA(2,7) \rightarrow R_1\right] = \frac{\langle WA(2,7) \rightarrow R_1 \rangle}{\langle \overline{WA}(2,7) \rightarrow R_1 \rangle} = 3$$

with no dependence upon the state of initial information about WA(2,7) before the appearance of R_1. On the other hand, it was natural to first introduce R_1 with the statement:

$$.75,\ R_1 \rightarrow WA(2,7)$$

because it focuses attention on the credibility of the fact which R_1 deals with. But if the fact WA(2,7) is already established at the .90 level, for example, we see how confusing such a method is for report characterization. This is why the normalizing assumption became necessary:

$$\langle WA(2,7) \rangle_{old} = 0.5$$

APPENDIX E

RULE CONSISTENCY

RULE CONSISTENCY

1. AN EXAMPLE OF A SELF-CONSISTENCY PROBLEM IN A SINGLE RULE

We use the following notation for a simple logical implication:

$$(.05,.70) \quad p \rightarrow q$$

where the two probability estimates are defined thus:

> $0.70 = \langle p \rightarrow q \rangle$ = probability estimate that the consequent q is true in the event that the antecedent p is true.
>
> $0.05 = \langle \neg p \rightarrow q \rangle$ = probability estimate that the consequent q is true in the event that the antecedent p is false.

So far there is no possibility of a failure in self-consistency, although the estimation of probability may be poor. For an antecedent of probability $\langle p \rangle$ = 0.75, *a Venn diagram with areas proportional to probability might appear thus:*

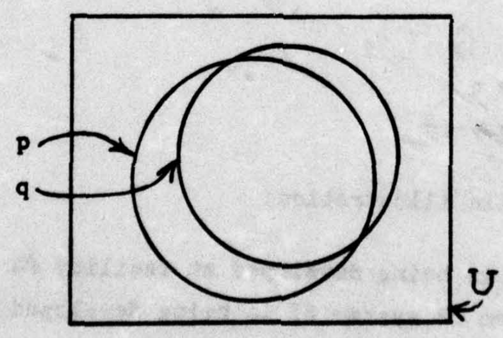

Assumed:

$\langle p \rangle = .75$
$\langle p \rightarrow q \rangle = .70$
$\langle \neg p \rightarrow q \rangle = .05$

With the aid of the Venn diagram interpretation it is fairly easy to gain conviction that if the .70 and .05 estimates are replaced by any other two probability numbers, the implication $p \rightarrow q$ is at least self consistent.

If, now, we consider the contrapositive of the original implication rule we might then obtain:

$$(.10,\ 0.80)\ \neg q \rightarrow \neg p$$
$$0.80 = \langle \neg q \rightarrow \neg p \rangle$$
$$0.10 = \langle q \rightarrow \neg p \rangle \qquad (\neg(\neg q) \equiv q)$$

We have defined the probability estimates in perfect analogy with the definitions for the original rule $p \rightarrow q$. We note, however, that the two probabilities are quite distinct from the original $\langle p \rightarrow q \rangle$, $\langle \neg p \rightarrow q \rangle$ probabilities. There is a relationship, of course, but it is not superficially obvious. We have also supposed that the .10, .80 estimates have been independently made.

It is natural to consider the contrapositive $\neg q \rightarrow \neg p$ because in simple logic it is fully equivalent to the original $p \rightarrow q$ statement. It is often used as a proof method, where to prove $p \rightarrow q$ we consider the reverse possibility $\neg q$. With the probabilities showing up as four independently estimated numbers, there is a definite possibility of a self inconsistency. This can be suspected from an examination of the Venn diagram, which has four areas which must add up to unity, so that there are only three degrees of freedom. We look therefore for a necessary relation amongst the four probability estimates, having assumed our rule in the two forms:

$$(.05,.70)\ p \rightarrow q$$
$$(.10,.80)\ \neg q \rightarrow \neg p$$

We can also assume the specific illustration:

$p \equiv$ system #7 is being developed at facility #4

$q \equiv$ a subsystem of system #7 is being developed at facility #4

AD-A036 069 AUERBACH ASSOCIATES INC PHILADELPHIA PA F/G 9/2
THE ANALYSIS OF CREDIBILITY AND CONSISTENCY IN INTELLIGENCE DAT--ETC(U)
DEC 76 J SABLE, R DICKSON F30602-75-C-0330
UNCLASSIFIED AAI-2329-TR-1 RADC-TR-76-392 NL

3 OF 3
AD
A036069

END
DATE
FILMED
3-77

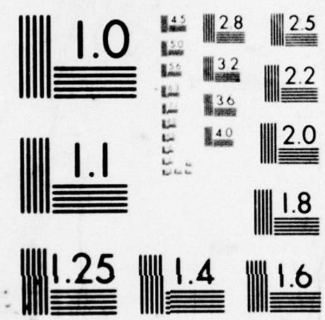

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963 A

We also assume that there are a total of 15 facilities. Superficially, the probability estimates appear plausible. For example, ($\langle \neg p \rightarrow q \rangle = .05$) if it is definitely known that system #7 is not being developed at facility #4, then there is a less than 1/15 chance that an associated subsystem is being developed there.

To obtain the necessary relationship among the four probability estimates we compute each of $\langle p \wedge q \rangle$ and $\langle p \vee q \rangle$ in two different ways:

$$\langle p \wedge q \rangle = \langle p \rangle \langle p \rightarrow q \rangle = \langle q \rangle \langle q \rightarrow p \rangle$$

$$\langle p \vee q \rangle = 1 - \langle \neg p \rangle \langle \neg p \rightarrow \neg q \rangle = 1 - \langle \neg q \rangle \langle \neg q \rightarrow \neg p \rangle$$

with the intention of eliminating $\langle p \rangle$, $\langle q \rangle$ amongst the equations to arrive at the relation amongst our four conditional probability estimates. The above are equivalent to:

$$\frac{\langle p \rangle}{\langle q \rangle} = \frac{\langle q \rightarrow p \rangle}{\langle p \rightarrow q \rangle} \qquad \text{(A)}$$

$$\frac{\langle \neg p \rangle}{\langle \neg q \rangle} = \frac{\langle \neg q \rightarrow \neg p \rangle}{\langle \neg p \rightarrow \neg q \rangle} \qquad \text{(B)}$$

Now we insert in the first relationship (A) the following:

$$\langle q \rangle = \langle p \rightarrow q \rangle \langle p \rangle + \langle \neg p \rightarrow q \rangle \langle \neg p \rangle$$

and obtain:

$$\langle p \rangle = \frac{1}{1 + \frac{\langle q \rightarrow \neg p \rangle \langle p \rightarrow q \rangle}{\langle \neg p \rightarrow q \rangle \langle q \rightarrow p \rangle}} \qquad (A^1)$$

Similarly, we insert in the (B) relationship:

$$\langle \neg q \rangle = \langle \neg p \rightarrow \neg q \rangle \langle \neg p \rangle + \langle p \rightarrow \neg q \rangle \langle p \rangle$$

and obtain:

$$\langle \neg p \rangle = \cfrac{1}{1 + \cfrac{\langle \neg q \rightarrow p \rangle \langle \neg p \rightarrow \neg q \rangle}{\langle p \rightarrow \neg q \rangle \langle \neg q \rightarrow \neg p \rangle}} \qquad (B^1)$$

Now we make use of $\langle p \rangle + \langle \neg p \rangle = 1$ in (A^1) and (B^1) and arrive at:

$$1 = \cfrac{1}{1 + \cfrac{\langle q \rightarrow \neg p \rangle \langle p \rightarrow q \rangle}{\langle \neg p \rightarrow q \rangle \langle q \rightarrow p \rangle}} + \cfrac{1}{1 + \cfrac{\langle \neg q \rightarrow p \rangle \langle \neg p \rightarrow \neg q \rangle}{\langle p \rightarrow \neg q \rangle \langle \neg q \rightarrow \neg p \rangle}}$$

We note that the four estimated probabilities are:

$$\langle p \rightarrow q \rangle = .70 \qquad \langle \neg p \rightarrow q \rangle = .05$$
$$\langle \neg q \rightarrow \neg p \rangle = .80 \qquad \langle q \rightarrow \neg p \rangle = .10$$

We also notice that the final symmetric relationship is what we have been looking for. It involves our probabilities directly (e.g., $\langle p \rightarrow q \rangle$) or through a co-probability (e.g., $\langle p \rightarrow \neg q \rangle$ where $\langle p \rightarrow \neg q \rangle + \langle p \rightarrow q \rangle = 1$). We are now in a position to check the degree of our probability estimation consistency:

$$\langle p \rangle + \langle \neg p \rangle = \cfrac{1}{1 + \cfrac{(.10)(.70)}{(.05)(1-.10)}} + \cfrac{1}{1 + \cfrac{(1-.80)(1-.05)}{(1-.70)(.80)}}$$

$$= \cfrac{1}{1 + \cfrac{14}{9}} + \cfrac{1}{1 + \cfrac{95}{120}}$$

$$= \frac{9}{23} + \frac{120}{215}$$

$$= .950$$

which we might identify as consistency in the measure of 95%

A person observing this might comment that perhaps $\langle \neg p \rightarrow q \rangle = .03$ might have been better than the .05 actually used. It is not easy to answer questions about improvement, but questions concerning consistency can be approached. The new check would appear thus:

$$\langle p \rangle + \langle \neg p \rangle = \frac{1}{1 + \frac{(.10)(.70)}{(.03)(1-.10)}} + \frac{1}{1 + \frac{(1-.80)(1-.03)}{(1-.70)(.80)}}$$

$$= \frac{1}{1 + \frac{70}{27}} + \frac{1}{1 + \frac{97}{120}}$$

$$= \frac{27}{97} + \frac{120}{217}$$

$$= .831$$

which indicates that the first estimation was more consistent. Since there has been no straddling, $\langle \neg p \rightarrow q \rangle = .057$ is apt to be quite nearly consistent with the other three probability estimates. The miss is by about 1%.

The literature speaks of people endowed with special capacities to estimate probability. We have here a specific test that can be employed to test the ability to estimate consistently, if not accurately.

2. EXAMPLES OF CONSISTENCY PROBLEMS BETWEEN A RULE AND OTHER INFORMATION

We comment, initially, that inconsistencies tend to appear in the form of disagreement between various estimates for the same probability. This is already apparent in the main body of the report where there have been cases of different portions of the data file leading to the same statement with very different probability estimates. This occurs because the statements used from the file may vary widely in credibility. In an even deeper sense, it is obvious that a direct clash, such as p and ¬p both being in the data file thus:

$$\left.\begin{array}{l} 1, p \\ 1, \neg p \end{array}\right\}$$

would be called a logical inconsistency in ordinary (deductive) logic; however in inductive logic the same situation:

$$\left.\begin{array}{l} 1,\ p \\ o,\ p \end{array}\right\}$$

is apt to be viewed as a difference in probability evaluation, or credibility.

Next we consider the consistency picture of a rule together with a fact statement:

$$\begin{array}{l} (.1,\ .9)\ \ p \Rightarrow q \\ \qquad\quad .95,\ q \end{array}$$

where we follow the notation of paragraphs 3.2. on.

$$\begin{array}{lllll} .9 & = & \langle p \rightarrow q \rangle & = & \text{prob. } (q/p) \\ .1 & = & \langle \neg p \rightarrow q \rangle & = & \text{prob. } (q/\neg p) \\ .95 & = & \langle q \rangle & = & \text{prob. } (q) \end{array}$$

Some consideration leads us to expect inconsistency directly. One statement asserts that assuming p (i.e., with its help) the truth probability of q is 90%. This appears to be at odds with the other statement that the truth probability of q is 95% without any help at all.

The inconsistency can be viewed formally by the estimates for $\langle q \rangle$. Working with the given rule alone:

$$\begin{aligned} \langle q \rangle &= \langle p \rightarrow q \rangle \langle p \rangle + \langle \neg p \rightarrow q \rangle \langle \neg p \rangle \\ &= (.9) \langle p \rangle = (.1) \langle \neg p \rangle \\ &= (.9) \langle p \rangle + (.1)\ (1 - \langle p \rangle) \\ &= (.8) \langle p \rangle + .1 \end{aligned}$$

so that the highest possible $\langle q \rangle$ is .9, whereas the data statement gives $\langle q \rangle = .95$. The view of the inconsistency cannot be shown in a Venn diagram with probability areas, unless we violate seriously some of the given information. We chose to violate the $\langle \neg p \rightarrow q \rangle = .1$ information. The Venn diagram, then, may appear thus:

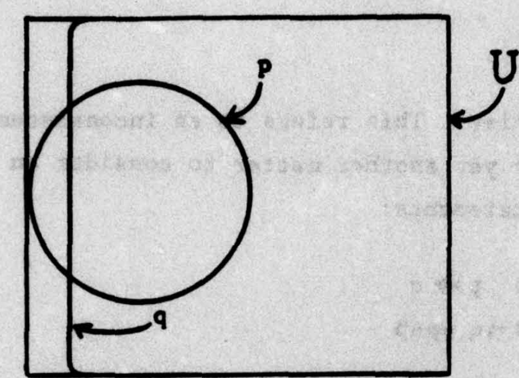

$$\langle q \rangle = .95$$
$$\langle p \rightarrow q \rangle = .90$$
$$\langle \neg p \rightarrow q \rangle = .96 \text{ (not .1)}$$

If we work with the Venn diagram attempt at (and distortion of) the information, and estimate $\langle q \rangle$ thus:

$$\langle q \rangle = \langle p \rightarrow q \rangle \langle p \rangle + \langle \neg p \rightarrow q \rangle \langle \neg p \rangle$$
$$.95 = .90 \langle p \rangle + .96 (1 - \langle p \rangle)$$
$$.06 \langle p \rangle = .01$$
$$\langle p \rangle = 1/6$$

Thus, given the large $\langle \neg p \rightarrow q \rangle = .96$ value it is possible to compute the value $\langle p \rangle$ from the assumed $\langle q \rangle$ value. This is the reverse direction of computing probability estimates as compared with our usual rule probability method. If this computation had been performed with the original $\langle \neg p \rightarrow q \rangle = .1$ assumption, then the inconsistency would have shown up as a $\langle p \rangle$ value greater than one.

Some general comments on the possibility of such an inconsistency may be helpful. If the .95, q statement is removed and replaced by the p statement with any probability estimate, e.g.:

.85, p

then no inconsistency with the given rule:

(.1, .9) $p \Rightarrow q$

can arise. Alternatively, if we consider a contrapositive form of the given rule, e.g.:

$$(.1, .9)\ \neg q \Rightarrow \neg p$$

with $.95$, q

then no inconsistency can arise. This refers to an inconsistency between the two statements above: it is yet another matter to consider an inconsistency between the following two statements:

$$(.1, .9)\ p \Rightarrow q$$
$$(.1, .9)\ \neg q \Rightarrow \neg p$$

Using the consistency test of Section 1 of this Appendix, we obtain a perfect consistency result. Independent of that test, we observe that:

$$\langle q\rangle = (.9)\langle p\rangle + (.1)\langle \neg p\rangle$$
$$\langle \neg p\rangle = (.9)\langle \neg q\rangle + (.1)\langle q\rangle$$

which are the relationships normally used in computing the probability estimate for the consequent of an implication. The result is that $\langle p\rangle = \langle q\rangle = 1/2$.

We proceed, now, with an example of inconsistency among essentially different rules (one rule not a contrapositive of another). Assume given:

$$(0, .9)\ p \Rightarrow q$$
$$(0, .9)\ q \Rightarrow r$$
$$(0, .9)\ r \Rightarrow p$$

An immediate expectation of inconsistency develops when one attempts a Venn diagram with probability areas. The first two rules must appear thus:

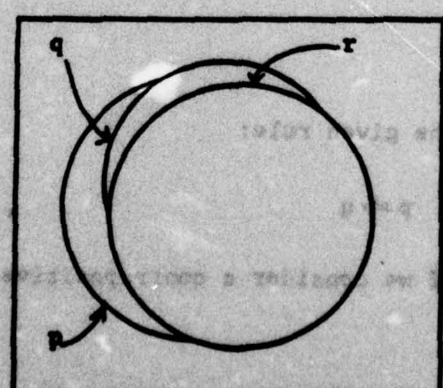

$$(0, .9)\ p \Rightarrow q$$
$$(0, .9)\ q \Rightarrow r$$

That is to say, the zero probability estimates ($\langle \neg p \rightarrow q \rangle = \langle \neg q \rightarrow r \rangle = 0$) mean that the q truth set is completely inside the p truth set. Likewise, the r truth set is completely inside the q truth set. Therefore, $\langle r \rightarrow p \rangle$ must equal unity instead of the .9 assumed in the third rule.

A more typical example of three such assumed rules follows:

$$(.1,\ .9)\quad p \rightarrow q$$
$$(.2,\ .9)\quad q \rightarrow r$$
$$(.3,\ .7)\quad r \rightarrow p$$

In this case, no inconsistency appears, as can be seen by solving the probability equations:

$$\langle q \rangle = .9\,\langle p \rangle + .1(1 - \langle p \rangle)$$
$$\langle r \rangle = .9\,\langle q \rangle + .2(1 - \langle q \rangle)$$
$$\langle p \rangle = .7\,\langle r \rangle + .3(1 - \langle r \rangle)$$

for which the solution is approximately:

$$\langle p \rangle \doteq .53$$
$$\langle q \rangle \doteq .52$$
$$\langle r \rangle \doteq .57$$

An attempt to do this with the earlier trio of rules would have led to a flat contradiction, and no possible solution. The insights of linear algebra in a $\langle p \rangle$, $\langle q \rangle$, $\langle r \rangle$ space appear valid. A Venn diagram with probability areas illustrating this last (and consistent) trio of rules follows:

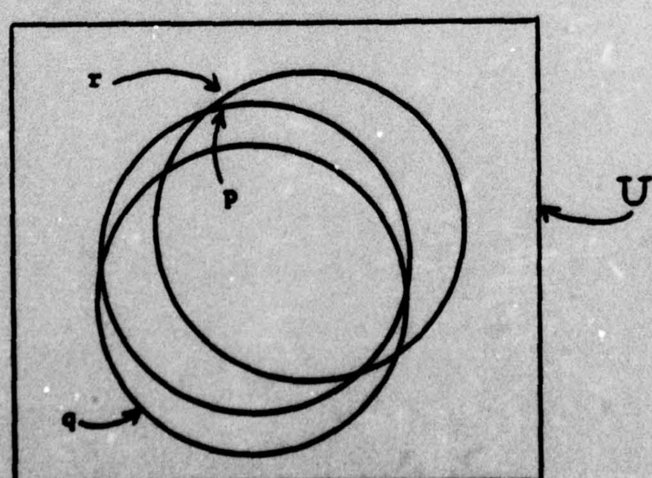

(The reverse of this page is blank)

APPENDIX F

A PAGING SCHEME FOR LISP

A PAGING SCHEME FOR LISP

The paging scheme will allow each user to perceive his own virtual space of 2^{17} words (131K) even though he is operating with a significantly smaller size of real core, probably on the order of 40K. This result can be accomplished by allowing the most needed portions of the user's virtual space to occupy whatever core space exists while the entire virtual space resides on drum. Whenever data not in core is needed by the user, that data is read from drum into whatever core space is not in use; if all core is in use, an equal sized, least needed section of data from the virtual space is removed from core to make room for the needed data. A table is used to keep track of where each section of virtual space is kept. In this scheme, the size of the section of virtual space is constant for all users; each such section is called a virtual block. The counterpart of the virtual block in core is the core block; similarly, the counterpart of the virtual block on drum is the drum block. While the size of virtual and core blocks is measured in words, the drum block size s is in sectors. A page is a virtual block allocated to the user.

The Page Table

Each block of virtual space has an entry in the Page Table. The n^{th} entry of the page table lists the current status of virtual block n. When the user references address @, the paging system divides @ by the page size to give a page number n and a page displacement d. The system looks at the n^{th} entry in the Page Table to find whether the page is core-resident. If so, the n^{th} entry on the Page Table will have the address a of the core block with which page n is identified. The system then reports to the user program the real address of @, namely a + d.

Page Fault

If page n is not core-resident, then it can be found on the drum. The user's space on the drum is identical to the virtual space, in the foll ing manner: virtual block n is identified with drum block n where drum block n is defined as residing on the s sectors starting with sector ns.

Now a place must be found for page n to occupy real core. First, the Available Core Block Pointer is consulted to see if a free core block is available. If so, the Available Core Block Pointer is changed to point to the next core block in the Available Core Block Stack and page n is read into core to occupy the available core block found. The Page Table is updated to reflect the change in page n's status.

If no core blocks are available, a page must be removed from core to make room for page n. Let us suppose we have determined page p has to go and so relinquishes its use of core block b. Page p is written out from core onto the drum and its Page Table entry reflects its change of location. Page n's Page Table entry is updated to show its acquisition of core block b and page n is read into block b in core.

Page Replacement Algorithm

The question still remains as to how the system decides which core-resident page must leave when a Page Fault occurs and no core blocks are available. First, the list of pages eligible for replacement is reduced by the Page Lock Feature; any page with its Page Lock "on" is not eligible for replacement.

The final decision is made using the Weight Counter stored in the Page Table for each core-resident page. Basically, every time a core-resident page is referenced its Weight Counter is incremented. So, the Weight Counter reflects the frequency of references to that page. When a page must be selected for replacement, the system scans the Weight Counters of all the (core-resident) pages to find the page with the lowest Weight Counter; this is the page to be replaced.

Now, the question is how much should a Weight Counter be incremented for a page reference? If the increment is a constant value, then the Weight Counter does not reflect the <u>current</u> need for the page in core. For example, suppose that over the last 10,000 page references, page 1 has been referenced 94 times while page 5 has been referenced 75 times. Let us further suppose

that all of page 1's references occurred over 5,000 page references ago (using page reference as a unit of time) while all of page 5's references occurred no more than 500 page references ago. It would seem reasonable that, although over the last 10,000 references page 1 has been referenced more frequently, page 5 has a better case for remaining in core because its references were more recent and, so, more closely reflect the current needs of the user. Thus, page references that occurred more recently should have higher increments or weights than those that occurred further in the past.

The function chosen as the basis for the page reference weight system is an exponential decay function of the form 2^{-ra}. Specifically, $w(r) = 2^{-r/a}$ is the weight of a page reference which occurred r page references ago where a is some positive constant. Clearly, function w satisfies our currency need that $w(x) > w(y)$ where $y > x \geq o$.

Using the results of simulation studies by Chu and Opderbeck[1], an estimate was made that our paging system could expect approximately 1 Page Fault for every 10,000 page references. Hence, a page reference occurring 10,000 page references ago should have a weight close to zero. W(10,000) was set equal to 1/16. So, our basis weight function was defined as $w(r) = 2^{-r/a}$ where $a = 2500$.

For computational reasons, the function w(r) has to be simulated using a step function s(x) in the following manner:

Let x be the x^{th} page reference after the most recent Page Fault.

$$
\begin{aligned}
\text{Then, } s(x) &= 1 \quad \text{for} \quad 0 < x \leq K \\
&= 2 \quad \text{for} \quad K < x \leq K + a \\
&= 4 \quad \text{for} \quad K + a < x \leq K + 2a \\
&= 2^n \quad \text{for} \quad K + (n-1)\,a < x \leq K + an
\end{aligned}
$$

where K is a function of a

A running count is kept of x. As each page reference occurs, x is increased by 1 and s(x) is added to the Weight Counter of the page referenced. When a

Page Fault occurs, a Factorization process is activated after the page to be replaced is found. The Factorization process merely divides each Weight Counter by the current $s(x)$. x is then reinitialized to zero and the paging system continues. In order to simulate $w(r)$, $s(x)$ is restricted so that when a Page Fault occurs at $x = an$ for $n \geq 1$, the sum of the weights of all page references from one to an inclusive is the same when calculated with $s(x)$ as with $w(r)$. That is, $\sum_{x=1}^{an} \frac{s(x)}{s(an)} = \int_{o}^{an} w(r)dr$ for all $n \geq 1$. $S(x)$ satisfies this condition when $K = a\ (2 - \frac{1}{\ln 2})$. With $a = 2500$, $K = 1393$. K is called Paging Parameter.K. A is called Paging Parameter A. X is the Page Reference Countdown.

In order to insure that a page just brought into core due to a Page Fault does not get replaced before it has a chance to accumulate page references, each page has its Weight Counter set to a Page Bonus value when it is brought into core. (Since locked pages are not eligible for replacement, references to such pages are ignored in the Weight Counter portion of the paging system.)

IMPLEMENTATION FIELDS

System Fields

Page Size = number of words/page; must be a power of 2

Core Size = number of core blocks

Paging Parameter K = number of page references with weight 1

Paging Parameter A = number of page references with weight 2^n where $n \geq 1$

Page Bonus = initial value of a page's Weight Counter when the page becomes core-resident

Virtual Size = number of blocks in virtual address space = maximum number of pages in initial memory

Drum Factor = number of sectors per page

User Fields (1 such field for each user)

Page Table - array of virtual block status entries

Page Reference Countdown - number of page references until weight is increased

Weight - number added to page's Weight Counter when page is referenced

Factor - number of bits each page's Weight Counter is shifted to the right during Factorization

Weight = 2^{Factor} at all times.

Page Table

The Page Table is an array of entries, one for each block in virtual memory. Each entry records the current status of a virtual block. The number of entries equals 2^{17}/Page Size.

Each Page Table entry will appear different according to the paging status of the virtual block

A) **Unallocated virtual block**

35	34 ... 18	17 ... 12	11 ... 0
1	1 ... 1	4	Available Virtual Block Pointer

Available Virtual Block Pointer = negative value of virtual block on Available Virtual Block Stack - if this virtual block is the last in the Available Virtual Block Stack, then this field contains the virtual block number associated with this entry (e.g., if virtual block 6 is the last virtual block in the Available Virtual Block Stack, its Available Virtual Block Pointer contains the value -6)

B) **Page that is out on drum**

35	34 ... 18	17 ... 12	11 ... 0
1	1 1 1	Type	1 1 ... 1

Type = type of LISP information stored on this page.

C) **Page that is core-resident**

35	34 ... 18	17 ... 12	11 ... 0
	Weight Counter	Type	Block Number

Page Lock (bit 35)

Page Lock = 1 if page must be core-resident
= 0 if page is not required to be core-resident

Weight Counter = weighted count of number of page references to this page since becoming core-resident

Block Number = number of block this page occupies in core; the real address of the core block can be obtained by shifting the block number n bits to the left where 2^n is the page size

Available Block Stack

A pushdown stack for each user is kept to record the blocks available to that user in core. The stack is structured as follows:

Available Core Block Pointer contains the address of a block available in core. If no blocks are available, the contents of the Available Core Block Pointer is zero.

In turn, the first word of the core block pointed to by the Available Core Block Pointer contains the address of the next available core block. If there are no more core blocks available, the first word contains zero.

In this manner, all available core blocks are listed on this single-threaded queue.

When the user needs an available core block, the user is given the first block on the stack and the Available Core Block Pointer is changed to point to the next block on the stack. When a core block is made available, that block's first word is set to point to the first block on the stack and the Available Core Block Pointer is changed to point to the newly available core block.

Available Page Stack

A pushdown stack for each user is kept to record the blocks available to that user in his virtual space. The stack is structured as follows:

The System Available Virtual Block Pointer contains the number of an available block in virtual space. If no such blocks are available, the content of the System Available Virtual Block Pointer is zero.

The Page Table entry of an available virtual block has an Available Virtual Block Pointer to record the next virtual block in the Stack. If no more blocks are available, the Available Virtual Block Pointer points to itself.

In this manner, all available virtual blocks are listed on this single-threaded queue.

REFERENCES

1. Chu, W.W. and H. Opderbeck, "Performance of Replacement Algorithms with Different Page Sizes," Computer, pp. 14-21, November 1974.

(The reverse of this page is blank)

APPENDIX G

THE LRU PAGING SCHEME FOR LISP

THE LRU PAGING SCHEME FOR LISP

Appendix F detailed a paging scheme modification to Univac 1100 LISP. The paging scheme included an algorithm to determine which page is to be replaced in the event of a Page Fault. This paper presents an alternate algorithm which can be used within the framework of the paging scheme previously presented. This algorithm is the so-called Least Recently Used (LRU) algorithm. A comparison of the two algorithms will be presented at the conclusion of this paper. All definitions and data structures from Appendix F are retained unless specifically noted.

The LRU Algorithm

In the LRU algorithm, the page number of each core-resident page is positioned in a pushdown stack according to how recently the page was referenced with the least recently referenced page at the bottom of the stack. So, whenever a page is referenced, its page number is removed from its current position in the stack and placed at the top. When a page must be replaced, the page chosen for replacement is the one at the bottom of the stack; the page brought into core is positioned at the top of the stack.

Changes in Implementation Structure

Using the LRU algorithm, the following system fields are not needed: Paging Parameter A, Paging Parameter K, Page Bonus. The following user fields are also not needed: Factor, Weight, Page Reference Countdown. The Weight Counter fields in the Page Table will be referred to as Reference Number fields. A single Next Reference Number field will be needed for each user. The system will also need a System Shift Factor field.

Implementation Algorithm

When a user begins, his Next Reference Number is initialized to zero. All Reference Number fields have the highest number they can hold (i.e., $2^{17} - 1$).

Whenever a page is referenced, the Next Reference Number value is placed in the Reference Number field of the appropriate Page Table entry. The Next Reference Number is then incremented by 1 and the paging system continues.

Whenever a page must be replaced, the Page Table is scanned to find the page with the lowest Reference Number. This is the page least recently used and is the one paged out; its Reference Number field is set to the highest value, $2^{17} - 1$ and the Page Lock bit is also set to 1. The page being paged in has its Reference Number set to the value of the Next Reference Number. The Next Reference Number is incremented by 1 and the paging system continues.

Whenever the Next Reference Number has the value $2^{17} - 1$ before being incremented, a Factorization process is set into motion to forestall the overflow situation. Each page's Reference Number field is shifted n bits to the right as is the Next Reference Number as well. The number n is the System Shift Factor. If n = 3, then Factorization will have to be performed only once for every 115,000 page references.

Comparison of Algorithms

In terms of the speed in the operation of the algorithms themselves, the LRU algorithm would seem to have the edge. In processing the ordinary page reference, the LRU algorithm does not have to perform as many arithmetic additions nor does it have to make as many comparison tests. When a page replacement is called for, the LRU algorithm has only to perform one comparison for each entry in the Page Table while the Weight Counter algorithm performs a comparison plus a shift for each entry. Both algorithms perform equally as fast in the overflow situation. However, the LRU algorithm will encounter the overflow situation once every 115,000 page references (when the System Shift Factor = 3). On the other hand, if Paging Parameter A = 2500, core size $\geq$ 80 pages, and the Page Fault Frequency (i.e., page faults per page reference) is greater than 1/25,000, overflow situations should almost never occur.

On the whole, the algorithms should give quite similar recommendations of pages for replacement. Still, there are differences. While the Weight Counter

algorithm emphasizes both currency and frequency, the LRU algorithm reflects only currency. Since currency is not balanced with frequency in the LRU, we can expect the LRU to more closely reflect the current needs of the user with no regard for any historical considerations of past need. The Weight Counter algorithm will be more moderate in its regard for the fad of the day. In some circumstances, though, the Weight Counter only reflects frequency. The Weight Counter algorithm because of its step function basis has the characteristic that over periods of high Page Fault Frequency on the order of 1/Paging Parameter K (e.g., 1/1,400 if Paging Parameter A = 2500) or more no regard at all is made for currency; the algorithm sees all page references with equal weight and, so, frequency rules supreme.

(The reverse of this page is blank)

METRIC SYSTEM

BASE UNITS:

Quantity	Unit	SI Symbol	Formula
length	metre	m	...
mass	kilogram	kg	...
time	second	s	...
electric current	ampere	A	...
thermodynamic temperature	kelvin	K	...
amount of substance	mole	mol	...
luminous intensity	candela	cd	...

SUPPLEMENTARY UNITS:

Quantity	Unit	SI Symbol	Formula
plane angle	radian	rad	...
solid angle	steradian	sr	...

DERIVED UNITS:

Quantity	Unit	SI Symbol	Formula
Acceleration	metre per second squared	...	m/s
activity (of a radioactive source)	disintegration per second	...	(disintegration)/s
angular acceleration	radian per second squared	...	rad/s
angular velocity	radian per second	...	rad/s
area	square metre	...	m
density	kilogram per cubic metre	...	kg/m
electric capacitance	farad	F	A-s/V
electrical conductance	siemens	S	A/V
electric field strength	volt per metre	...	V/m
electric inductance	henry	H	V-s/A
electric potential difference	volt	V	W/A
electric resistance	ohm		V/A
electromotive force	volt	V	W/A
energy	joule	J	N-m
entropy	joule per kelvin	...	J/K
force	newton	N	kg-m/s
frequency	hertz	Hz	(cycle)/s
illuminance	lux	lx	lm/m
luminance	candela per square metre	...	cd/m
luminous flux	lumen	lm	cd-sr
magnetic field strength	ampere per metre	...	A/m
magnetic flux	weber	Wb	V-s
magnetic flux density	tesla	T	Wb/m
magnetomotive force	ampere	A	...
power	watt	W	J/s
pressure	pascal	Pa	N/m
quantity of electricity	coulomb	C	A-s
quantity of heat	joule	J	N-m
radiant intensity	watt per steradian	...	W/sr
specific heat	joule per kilogram-kelvin	...	J/kg-K
stress	pascal	Pa	N/m
thermal conductivity	watt per metre-kelvin	...	W/m-K
velocity	metre per second	...	m/s
viscosity, dynamic	pascal-second	...	Pa-s
viscosity, kinematic	square metre per second	...	m/s
voltage	volt	V	W/A
volume	cubic metre	...	m
wavenumber	reciprocal metre	...	(wave)/m
work	joule	J	N-m

SI PREFIXES:

Multiplication Factors	Prefix	SI Symbol
1 000 000 000 000 = 10^{12}	tera	T
1 000 000 000 = 10^{9}	giga	G
1 000 000 = 10^{6}	mega	M
1 000 = 10^{3}	kilo	k
100 = 10^{2}	hecto*	h
10 = 10^{1}	deka*	da
0.1 = 10^{-1}	deci*	d
0.01 = 10^{-2}	centi*	c
0.001 = 10^{-3}	milli	m
0.000 001 = 10^{-6}	micro	µ
0.000 000 001 = 10^{-9}	nano	n
0.000 000 000 001 = 10^{-12}	pico	p
0.000 000 000 000 001 = 10^{-15}	femto	f
0.000 000 000 000 000 001 = 10^{-18}	atto	a

* To be avoided where possible.

MISSION
of
Rome Air Development Center

RADC plans and conducts research, exploratory and advanced development programs in command, control, and communications (C^3) activities, and in the C^3 areas of information sciences and intelligence. The principal technical mission areas are communications, electromagnetic guidance and control, surveillance of ground and aerospace objects, intelligence data collection and handling, information system technology, ionospheric propagation, solid state sciences, microwave physics and electronic reliability, maintainability and compatibility.

AMERICAN REVOLUTION BICENTENNIAL 1776-1976

☆U.S. GOVERNMENT PRINTING OFFICE: 1977-714-025/94